

End-to-End Encrypted Zoom Meetings: Proving Security and Strengthening Liveness* **

Yevgeniy Dodis^{1***}, Daniel Jost^{1***}, Balachandar Kesavan², and Antonio Marcedone²

¹ New York University {dodis, daniel.jost}@cs.nyu.edu

² Zoom Video Communications {balachandar.kesavan, antonio.marcedone}@zoom.us

Abstract. In May 2020, Zoom Video Communications, Inc. (Zoom) announced a multi-step plan to comprehensively support end-to-end encrypted (E2EE) group video calls and subsequently rolled out basic E2EE support to customers in October 2020. In this work we provide the first formal security analysis of Zoom’s E2EE protocol, and also lay foundation to the general problem of E2EE group video communication.

We observe that the vast security literature analyzing asynchronous messaging does not translate well to synchronous video calls. Namely, while strong forms of forward secrecy and post compromise security are less important for (typically short-lived) video calls, various *liveness* properties become crucial. For example, mandating that participants quickly learn of updates to the meeting roster and key, media streams being displayed are recent, and banned participants promptly lose any access to the meeting. Our main results are as follows:

1. Propose a new notion of *leader-based continuous group key agreement with liveness*, which accurately captures the E2EE properties specific to the synchronous communication scenario.
2. Prove security of the core of Zoom’s E2EE meetings protocol in the above well-defined model.
3. Propose ways to strengthen Zoom’s liveness properties by simple modifications to the original protocol, which subsequently influenced updates implemented in production.

* This is the full version of the article with the same title published in the Proceedings of EUROCRYPT 2023, Springer, © IACR 2023.

** In this version, we clarified that the protocol deployed by Zoom since v5.13 of the client has some differences with respect to the abstraction presented in the paper, and corrected some inaccuracies in the theorem statements and protocol description with respect to the original proceedings version that do not fundamentally change our results.

*** Research conducted while contracting for Zoom

Table of Contents

1	Introduction	3
1.1	Contributions	4
1.2	Related Work	5
1.3	Organization	6
2	Continuous Multi-Recipient KEM	6
2.1	Syntax	7
2.2	PKI	8
2.3	Security Definition	8
2.4	Zoom’s Scheme	9
3	Leader-based GCKA with Liveness	12
3.1	Syntax	12
3.2	Security Definition	13
3.3	Zoom’s Scheme	14
4	Improved Liveness	18
4.1	Limitations of Zoom’s Protocol	18
4.2	Additional Interaction	19
4.3	Leveraging Clock Synchronicity	21
5	Meeting Stream Security	23
6	Conclusions	24
7	Acknowledgements	24
A	Preliminaries	27
A.1	Notation	27
A.2	Pseudocode and Games	27
A.3	Cryptographic Primitives	27
A.4	Gap Diffie Hellman	29
B	Zoom’s PKI	29
C	Details on cmKEM	30
C.1	The PKI	30
C.2	Correctness	30
C.3	Security	30
C.4	Zoom’s cmKEM Scheme	34
D	Details on LL-CGKA	38
D.1	Security	38
D.2	The protocol	42
D.3	Proof of Theorem 2	43
D.4	Correctness	47
E	Details on Improved Liveness	54
E.1	Additional Interaction: Proof of Theorem 3	54
E.2	Additional Interaction: Correctness	55
E.3	Leveraging Synchronicity: Proof of Theorem 4	56
E.4	Leveraging Synchronicity: Correctness	60

1 Introduction

Group video communication tools have gained immense popularity both in personal and professional settings. They were instrumental in bringing people closer together at a time when travel and in-person interaction were severely limited by the COVID-19 pandemic. Zoom Video Communications, Inc. (Zoom) is one of the leading providers of video communications with millions of active users, and aims to distinguish itself not just in ease-of-use and richness of features, but also by offering strong security and privacy capabilities.

Historically, Zoom meetings have been encrypted in transit between the clients and the Zoom servers. This allows Zoom to offer features that require the server to access meeting streams, such as live transcription and the ability to join a meeting by dialing a phone number through the telephony network. In May 2020, Zoom announced a multistep plan to comprehensively support end-to-end (E2E) encrypted group video calls [47] and rolled out basic E2EE support to the public in October 2020 [33]. E2EE protects the privacy of attendees against any compromise to Zoom’s infrastructure/keys.

Zoom has also published a whitepaper [11] describing its protocol, design goals, and methodology for E2EE meetings. The whitepaper explains how the protocol is run as part of a group call and provides intuition on the threat model and security. Subsequent academic work has performed an initial analysis of the protocol [29], emphasizing a number of potential attacks at the boundary of the threat model outlined in the whitepaper. However, this security analysis is far from comprehensive and does not include any formal security definitions or theorems.

Group video calls. E2EE group video calls have not gained any major scrutiny from the academic community. This stands in stark contrast to related fields such as secure text messaging, where the ubiquitously used Signal protocol [35] has received significant attention [2,17,10]. For secure group messaging, the Internet Engineering Task Force (IETF) has even launched the Messaging Layer Security (MLS) working group [8] with mutual support from industry and academia, resulting in a number of analyses [3,4,5].

A defining feature of any group video call that distinguishes it from the asynchronous nature of text messaging is that video calls happen in real-time with all participants online at the same time. This suggests that protocols could achieve strong *liveness properties* generally deemed to be intrinsically unattainable in messaging. First, an attacker should not be able to arbitrarily delay communication. For example, if Alice sends a video stream at time t , then Bob should not accept it at a time significantly later than t . Depending on the type of content, such delays may pose a significant threat; for instance, if the message is an instruction to buy or sell a certain stock, then the ability to delay it might allow an attacker to front-run the transaction. Second, if the meeting host decides on a certain management action, such as adding or removing parties, then an attacker must not be able to delay or prevent those decisions from taking effect. These liveness properties are new and not demanded in the (asynchronous) group messaging setting, in which the network attacker can simply pretend that the initiating party is offline, without any of the other parties being able to detect the attack.

Goals of this work. In this work we aim to analyze the core of Zoom’s E2EE meetings protocol¹. We follow the approach successfully used to analyze (group) messaging protocols and single out the key agreement using the abstraction of a so-called *continuous group key agreement (CGKA)* protocol [3], albeit with weaker forward secrecy (FS) and post-compromise security (PCS) properties than for secure messaging, as explained below. The CGKA abstraction establishes a sequence of shared symmetric keys, accounting for the need to re-key when parties join or leave the meeting (even without strong FS/PCS). The current key — known exactly to the current members of the meeting — can then be used with authenticated encryption with associated data (AEAD) to achieve secure communication.

To provide the first formal security analysis of Zoom’s E2EE protocol, our main objectives are, thus, to:

1. Propose a CGKA definition that takes Zoom’s unique aspects into account and captures the liveness properties made possible by the online assumption.

¹ We analyze the Zoom E2EE meetings protocol as deployed in the Zoom client version 5.12. In this paper, we refer to this version as the *current* protocol/scheme.

2. Provide an analysis of the core of Zoom’s E2EE protocol in the above well-defined model.

To the best of our knowledge, Zoom is the only E2EE group video protocol that aims to provide stringent liveness properties. We believe our work is the first in the realm of CGKA to formalize and analyze such assurances. As part of this process, we observed that Zoom’s liveness assurances could be strengthened and thus, we set out to:

3. Propose tangible strengthenings to Zoom’s liveness properties, via two simple modifications to the protocol which offer different tradeoffs between efficiency and security. Zoom has evaluated these modifications and successfully deployed a variant of one of them in production (in version 5.13 of the Zoom client).

1.1 Contributions

Definition. We formally define a *leader-based continuous group key agreement with liveness* (LL-CGKA), which encompasses all the desired security properties of Zoom’s core E2EE meetings protocol in a single security game¹. In general terms, an LL-CGKA protocol requires the following properties:

- At each stage of the meeting, the shared symmetric key is only known to the set of current participants as decided by the current meeting host.
- All participants have a consistent view of the set of current meeting participants (as displayed in the UI) as well as of the key.
- Changes to the group, decided by the meeting host, are applied within a bounded (and short) amount of time; otherwise, participants drop out of the meeting.

Attacker model. We consider a powerful adversary that has control over the evolution of the group, fully controls the network and Zoom’s server infrastructure, and can passively corrupt any parties, thereby obtaining their current state. We remark, however, that most of our guarantees hold only when the current meeting leader and participants execute the protocol honestly, and any active attackers previously in the meeting have been removed.

FS and PCS. Due to the short-lived nature of video calls, our CGKA notion differs from those in realm of secure messaging by requiring neither strong forward secrecy nor post-compromise security guarantees within a single meeting. An attacker compromising a party’s state in an ongoing meeting may learn both past and future content of said meeting. We do, however, require the following properties: first, corrupting a party must not reveal any of the meeting’s content before the party has been admitted or after it has been removed by the meeting host. Second, compromising a party after a meeting has ended must not compromise the meeting in any form (weak FS). Third, even if a party’s long-term secret have been leaked, this party can still securely join meetings, maintaining confidentiality as long as the adversary does not act as an active meddler-in-the-middle.

Modularization. One of the contributions of this work is to distill out basic building blocks of Zoom’s protocol, which could be instantiated differently in pursuit of improved efficiency or, e.g. to achieve post-quantum security. To this end, we consider the intermediate *continuous multi-recipient key encapsulation* (*cmKEM*) abstraction from which we then build the aforementioned LL-CGKA notion. Put simply, the former naturally captures that in Zoom’s protocol a designated party (the meeting host) chooses the symmetric key and distributes it to all the meeting participants. The latter abstraction then models the core of Zoom’s E2EE meetings protocol, including the unique liveness properties.

Finally, we discuss how Zoom’s overall protocol is built on top of the LL-CGKA protocol, considering audio and video encryption. In particular, we relate the respective confidentiality, authenticity and liveness assurances to those of the LL-CGKA notion.

We remark that the above modularization follows Zoom’s whitepaper [11] version 4.0, with the cmKEM notion roughly corresponding to Sections 7.6.2 – 7.6.6, the LL-CGKA notion to Section 7.6.7, and video stream encryption discussed in Sections 7.2 and 7.11, among others.

Liveness. One of the main novelties of Zoom’s E2EE protocol is its focus on liveness properties. They assure that whenever the host adds or removes a participant, the action cannot be withheld by an adversary for any extended period of time. That is, if for instance the host removes a member from the group, such as when removing a candidate at the end of an interview so that the hiring panel can reach a decision, that member must no longer be able to decrypt meeting contents even if they manage to compromise Zoom’s cloud infrastructure or exert significant control over the network.

In this work, we present a simple time-based model that allows us to formalize and analyze those liveness properties. Our model balances simplicity and generality by assuming that parties have access to local clocks that all run at the same speed, but are otherwise not assumed to be synchronized. We then formalize liveness as follows: whenever a participant is in a given state at time t , then the meeting host has been in the same state recently, i.e., at some time $t' \geq t - \Delta$ where Δ is some protocol-dependent liveness slack. Turned around, whenever the host moves on to a new state (e.g., by changing the group roster) then all participants must also move on within time Δ (or else drop from the meeting).

While the protocol we analyze¹ achieves good liveness properties, these assurances degrade in the number of host changes. As part of this paper we propose two potential improvements. First, we propose a modification that strictly improves on the liveness and yields bounds independent of the number of host changes. This comes at the cost of increased communication by making the protocol more interactive. As an alternative, we propose a strengthening that does not incur any communication overhead and improves on Zoom’s properties if parties have well-synchronized clocks; we believe this to be the common case for modern devices. After testing, Zoom implemented a variant of the first option, which was deployed in version 5.13 of the Zoom client.

1.2 Related Work

We have already commented above on the relationship of this work to the areas of secure messaging.

Group video calls. There are numerous solutions for group video calls. The vast majority offers transport layer encryption, with some of them [7,16,11,44,45] offering E2EE group calls, and others offering this feature only for two-party calls [27,35]. While some of the solutions do offer intuitive security descriptions in the form of a whitepaper, such as Wire [45], Cisco [16], and WhatsApp [44], to the best of our knowledge only Cisco WebEx enjoys formal security claims, as it is directly built on top of the IETF MLS draft [8,3,4,5].

Liveness. The terms liveness, liveliness, and aliveness are frequently used to describe various *authentication properties* of key agreement protocols, e.g., in [34] (and many subsequent works). Those properties, roughly speaking, guarantee that if one party completes a run of the protocol, then its peer at some point also has run the same protocol. (Slightly stronger variants exist.) As such most of those definitions not only have no direct relation to *physical time* but also are typically not enforced on an ongoing basis, contrary to our liveness definition. Further, in the context of E2EE group messaging, some work previously used the liveness as synonymous to correctness [40] — with no direct relation to actions having to occur in a timely manner.

However, using timing is not new in the design and analysis of cryptographic protocols. Some such works (e.g., [38,23,30]) use timing assumptions to improve efficiency (or overcome impossibility results) for problems which do not inherently require timing assumptions. Other works (e.g., [22,41,9,12]) use various forms of “moderately hard function” to achieve different cryptographic properties which critically rely on the notion of time. The type of liveness used in this work is much more closely related to more traditional distributed computing literature (e.g., [21,24]) on consensus and, more recently, blockchain protocols (e.g., [26,37]). However, the existence of a unique meeting leader, coupled with the online assumption, makes Zoom’s protocol (and our security model) much more lightweight. Finally, the use of heartbeats to ensure liveness is similar to the heartbeat extension of the TLS protocols [42].

Related Notions. The cmKEM notion is an extension of multi-recipient Key Encapsulation (mKEM) [43,39,46] to the setting of dynamically changing groups. Zoom’s scheme is based around the authenticated

public-key encryption² scheme from the `libsodium` library [20]. It is very similar to one of the early authenticated public-key encryption schemes formally analyzed by An [6] (and simpler than the recently analyzed HPKE standard [1]).

The LL-CGKA notion is further related to Dynamic Group Key Agreement with an extensive body of literature, notable examples including [14,28,31]. Similar to CGKA, the Dynamic GKA notion supports changes to group membership during a session and, in fact, in terms of FS and PCS guarantees those notions resemble our LL-CGKA notion more closely than most prior CGKA variants. In contrast to CGKA, Dynamic GKA schemes are designed for an interactive setting, i.e., typically require all parties to contribute to any one operation via interactive rounds, and / or rely on a trusted group manager. (In contrast to LL-CGKA the group manager is, however, static and cannot be replaced mid-session.) Further, we note that while group video calls in principle can tolerate interactive protocols, such as [31], requiring several parties to contribute to each operation can be nevertheless problematic, as for example parties can unexpectedly drop out. Furthermore, we believe this simplifies extending our notion for a more advanced group video call protocol, compared to a Dynamic GKA based one. Closely related to Dynamic GKA are further Multicast Encryption, e.g., [36], and the line of work on Logical Key Hierarchies, e.g., [15].

Another related notion to both cmKEM and LL-CGKA is Multi-Stage Authenticated Key Exchange [25]. Several variants, each with slightly different guarantees, have been considered and the notion has, e.g., been used to analyze the Double Ratchet protocol [18]. In contrast to CGKA, Multi-Stage AKE has exclusively been applied to the two-party setting.

1.3 Organization

In Section 2, we formalize the cmKEM notion, in which one participant in a meeting generates and distributes key material to the remaining participants, and describe Zoom’s implementation of this building block. In Section 3, we define the LL-CGKA notion, which strengthens the cmKEM definition with properties such as liveness and group consistency, and present an abstraction of Zoom’s LL-CGKA protocol¹. In Section 4, we point out limitations in Zoom protocol’s liveness guarantees, and propose two improved protocols, one of which influenced an update to Zoom’s protocol. Finally, in Section 5, we comment on the confidentiality, authenticity, and liveness guarantees of the meeting stream contents itself.

Appendix A describes the syntax, primitives, and cryptographic assumptions used in this paper. Appendix B describes the PKI used by Zoom for long-term key verification. Appendices C, D and E provide more details and formal security proofs about the cmKEM notion, the LL-CGKA notion, and our proposed liveness improvements, respectively.

2 Continuous Multi-Recipient KEM

Zoom’s protocol works by having a designated party distribute shared symmetric key material to all the participants upon each change to the group. We abstract this as a *Continuous Multi-Recipient Key Encapsulation* (cmKEM) scheme that allows the designated party to encapsulate a stream of shared symmetric keys to a dynamically evolving set of recipients. This results in a sequence of independent and uniformly random *keys*, each only known to the authorized parties. We number the states (i.e., keys) using two counters: the *epoch* and a sub-epoch called *period*.³

In the following, we call the designated party *leader*.⁴ We ignore policy aspects and assume that the leader is told whom to add or remove.

² Authenticated public-key encryption schemes are often also referred to as *signcryption schemes*. The latter term is however more commonly used to denote schemes satisfying insider security rather than outsider security, as achieved by `libsodium`’s scheme.

³ Looking ahead, rotating the period instead of the full epoch during group additions is more efficient. Zoom’s protocol currently does not take advantage of period rotations, but we capture and analyze this option since it is being considered as a future optimization.

⁴ Typically the leader coincides with the meeting host, but if, e.g., the host is on a low-bandwidth connection, these concepts can be decoupled.

The cmKEM notion distinguishes long-term identities and *ephemeral users*. Each long-term identity id is assumed to have an associated public key ipk . A party id can then create one or more ephemeral users, identified by uid , each linked to a specific *meeting*. That is, each meeting will consist of a group of ephemeral $uids$ that just exist for the duration of that meeting. Roughly speaking, in Zoom, each long-term identity id corresponds to a device; if a user logs into multiple devices, each will have its own long-term key material. Note that a device can be part of the same meeting under different ephemeral identities over time, e.g., after leaving the meeting and then rejoining it.

To cope with the leader suddenly losing connection, leader switches are initiated by the (untrusted) server without any hand-off. As a result, a user uid can be asked at any point of time to become the new leader of a meeting, with any given set of participants, as long as they are associated with the same meeting. To simplify notation, we introduce the notion of a *session* that denotes a segment of meeting between leader changes.

2.1 Syntax

For simplicity, we define the clients' cmKEM algorithms to be non-interactive, making all the interaction explicit by having multiple algorithms. User algorithms moreover have implicit access to a PKI as described in the next section. The server aids the protocol execution by performing explicit message routing.

Definition 1. A cmKEM scheme consists of the following algorithms. For ease of presentation, the client state st is assumed to expose the current key $st.k \in \mathcal{K}$, epoch $st.e$, and period $st.p$, while the server state pub is assumed to expose a mapping $pub.E[\cdot]$ from meetings to the meeting's current epoch.

User management:

- $(st, uid, sig) \leftarrow \text{CreateUser}(id, meetingId)$ creates an ephemeral user belonging to id and the meeting $meetingId$. It outputs the initial state st , the user's identity uid , and credential sig binding uid to id .
- $(id, ipk) \leftarrow \text{Identity}(uid)$ and $meetingId \leftarrow \text{Meeting}(uid)$ deterministically compute uid 's long-term information, and associated meeting respectively.

Session management:

- $(st', M) \leftarrow \text{StartSession}(st, \{(uid_i, ad_i, sig_i)\}_{i \in [n]}, e')$ instructs the user to start a new session with the given members. For each member, credential sig_i as well as associated data ad_i (which needs to match with the user's respective value when joining) are provided. The algorithm takes an optional argument e' indicating the starting epoch for the session. The welcome message M is to be distributed to the other group members by the server.
- $st' \leftarrow \text{JoinSession}(st, uid_{lead}, sig_{lead}, m, ad)$ makes the user join the leader's session using their share m of the welcome message.

Group and management (leader):

- $(st', M) \leftarrow \text{Add}(st, \{(uid_i, ad_i, sig_i)\}_{i \in [n]}, newEpoch)$ adds the users uid_1 to uid_n to the group. The boolean flag $newEpoch$ indicates whether this action should create a new epoch or period. This algorithm can also be called without any uid tuple as input in order to introduce a new key without changing the set of participants.
- $(st', M) \leftarrow \text{Remove}(st, \{uid_i\}_{i \in [n]})$ removes the users uid_1 to uid_n from the group.

Message processing (non-leaders):

- $st' \leftarrow \text{Process}(st, m)$ lets a participant advance to the next epoch or period.

Message passing (server):

- $pub \leftarrow \text{InitSplitState}()$ generates an initial public server state.
- $(pub, \{(uid_i, m_i)\}_{i \in [n]}) \leftarrow \text{Split}(pub, M)$ deterministically splits M into shares m_i for each recipient.

Correctness is formally defined in Appendix C.2.

2.2 PKI

Each ephemeral user identity uid is bound to its long-term identity id via a signed credential. To this end, id has a long-term signing key isk . In order to prevent meddler-in-the-middle (MITM) attacks, other parties must authenticate the respective long-term public key ipk . For the sake of our analysis, we assume a simple (long-term) public-key infrastructure (PKI). The PKI provides to each long-term identity id their respective private signing key isk while allowing all other users to verify that the respective public verification key ipk belongs to id .

Zoom currently does not have any such PKI but relies on the host reading out a *meeting leader security code* — a digest of ipk — that all participants then compare to ensure they have the host’s correct key. Authentication crucially depends on the leader visually recognizing participants and vice versa. Formalizing the exact guarantees given by this process is outside the scope of this work — specifically because the authenticity is only established during and not before a meeting, and because it relies on non-cryptographic assumptions such as the host recognizing participants’ faces.

In the future, Zoom plans to build a PKI based on key transparency and external identity providers; analysis thereof is left for future work. We refer to Appendix B for a more in-depth discussion on how Zoom currently verifies public keys, as well as their ongoing efforts for improving user authentication.

2.3 Security Definition

The security notion for the cmKEM primitive encompasses all the desired security properties in a single game. We next describe its the high-level workings, with the full formal definition presented in Appendix C.3.

Game overview. The attacker has full control over the evolution of the group and the network. We now sketch the various oracles the adversary may call. First, the adversary can *create a user* for a provided long-term identity id and meeting $meetingId$. The game ensures that the generated user uid is unique.

The adversary can then instruct uid_{lead} to *start a session* for a provided list of participants and their respective credentials. Afterwards, they can instruct a user uid to *join uid_{lead} ’s session* using a welcome message m of the adversary’s choice. The leader can also be instructed to *add or remove members*. In the former case, it is up to the adversary to specify whether this should initiate a new epoch or period. Finally, the adversary can get a participant uid to *process an arbitrary message m* , though the protocol might of course reject malicious messages.

The game ensures that additions and removals only succeed if the adversary does not try to add existing members or remove nonexistent ones. Additionally, the leader must not remove themselves from the group; instead, another party could assume the role of the leader and exclude the old leader from the group. The game keeps track of, for each leader’s epoch and period, the leader’s view of the session state, which consists of the meeting key and participant roster. Throughout the execution, the game then ensures consistency of the parties’ view with their leader’s respective view, which we discuss below.

The attacker can passively corrupt long-term identities, which reveals (a) the secret states of all still active associated ephemeral identities and (b) the long-term identity’s signing key from the PKI.

Key confidentiality. The adversary must not be able to distinguish the keys produced by the cmKEM scheme from random ones. To this end, the adversary may try to guess a bit b by challenging a state’s key (identified by the leader, epoch, and period) to either receive the real key (if $b = 0$) or a uniform random one (if $b = 1$). Additionally, the game allows the adversary to instead request the actual key, irrespective of the bit b , which may be useful since the adversary cannot win if it corrupts any party who knows any of the challenged keys.

The game needs to rule out trivial wins stemming from the adversary being able to compute certain keys themselves after passively corrupting parties. Since Zoom’s scheme neither encompasses forward secrecy (FS) nor post-compromise security (PCS) within a session, this has to be reflected in our notion. In short, corrupting a user potentially reveals the key for all epochs and periods where he has been a member of a given session. However, keys must remain secure in the following situations:

- A user must never know keys from before being added to, or after having been removed from the group. Hence, the confidentiality of those keys must not be affected by compromising the given user.
- Corrupting a device after a session has ended, i.e., after the respective ephemeral identity has been deleted, must not affect the session’s confidentiality.⁵
- Corrupting a long-term identity `id` (and thus learned `isk`) must not affect the security of future sessions involving an honestly generated ephemeral identity `uid` for `id`. (The adversary might of course impersonate `id` by creating a valid ephemeral user `uid'` instead, which would compromise the session’s security.)

Consistency properties. Parties must agree on the key for each epoch and period within a given session. That is, no two parties should ever output conflicting keys, unless after an active attack in which the adversary uses either the leader’s or the receiving party’s leaked state to tamper with the messages.⁶ Consistency, moreover, takes into account at which point in time parties can reach a given state. Our notion distinguishes between epochs and periods, among others, due to those properties differing. Participants must only move to an epoch once their leader has arrived there, while for periods, we allow participants to run ahead and thus reach periods that formally are not supposed to exist. (Still, parties must agree on the keys for those spurious periods.)

Finally, consistency must hold even if the adversary tampers with, reorders, or replaces messages — as long as the involved parties are honest. Due to the leader-based nature of the cmKEM primitive, a malicious leader however could always break consistency by simply sending inconsistent messages to the respective parties. To formalize outsider security, we thus simply deem attacks enabled by corrupting one of the involved parties trivial and no longer enforce consistency properties for a user `uid` once either `uid` or their leader `uidlead` has been corrupted.

Member authentication. For many operations such as adding users to an existing session or instructing a user to join another session, the adversary is allowed to provide the respective user identifiers. Our security notion ensures that the adversary cannot impersonate long-term identities unless they have been corrupted, i.e., the adversary cannot inject an ephemeral user `uid` unless the associated long-term identity `id` has been corrupted.

2.4 Zoom’s Scheme

Zoom’s cmKEM scheme uses point-to-point encryption to communicate fresh keys to the participants — i.e., it does not leverage any efficiency gains from sending the same message to multiple recipients. It is based on Diffie-Hellman key exchange over a cyclic group $\mathbb{G} = \langle g \rangle$ with a fixed generator g . The identifier `uid` mainly consists of a Diffie-Hellman public key `upk` $\in \mathbb{G}$ alongside the contextual data of the meeting identifier, the user’s long-term identity `id`, the user’s long-term public key `ipk`, and a signature under the user’s long-term signing key `isk` binding it all together. The respective secret key `usk` $:= \text{DLog}_g(\text{upk})$ is stored as part of the protocol’s state. See Fig. 1 for a formal description of the scheme.

For each epoch, the leader samples a new *seed*, from which the sequence of period keys are derived by iteratively applying a PRG to derive the key and seed for the next period of that epoch³. (Observe that this construction is forward-secure.) When removing parties, the leader initiates the next epoch and communicates the new seed to all remaining participants as described below. They then all derive the first key and the seed for the second key using the PRG. Analogously, to add participants with `newEpoch = true`, the leader communicates the seed to all participants. More efficiently, however, when adding participants with `newEpoch = false`, the leader only sends the seed for the next key to the freshly joined parties and instructs the others to just ratchet forward.

To send a seed to a party, the leader first derives for each recipient a shared symmetric key from a Diffie-Hellman element computed from its own secret key `usk` and the recipient’s public key `upk'`. The scheme

⁵ As in TLS, we require FS on the granularity of sessions.

⁶ This formalizes an outsider notion actually achieved by Zoom. Stronger protocols could tolerate leaking the recipient’s state.

Protocol cmKEM Client Protocol

User management

Algorithm: CreateUser(id, meetingId)
 $usk \leftarrow \$_\{0, 1, \dots, |G| - 1\}$; $upk \leftarrow g^{usk}$
 $(isk, ipk) \leftarrow PKI.get-sk(id)$ // id's long-term keys
 $me \leftarrow (meetingId, id, ipk, upk)$
 $sig \leftarrow Sig.Sign(isk, 'EncryptionKeyAnnouncement', me)$
 $K[\cdot], uid_{lead}, G, k, e, p, seed \leftarrow \perp$
return (me, sig)

Algorithm: Meeting(uid)
parse (meetingId, id, ipk, upk) $\leftarrow uid$
return meetingId

Algorithm: Identity(uid)
parse (meetingId, id, ipk, upk) $\leftarrow uid$
return (id, ipk)

Session management

Algorithm: StartSession($\{(uid_i, ad_i, sig_i)\}_{i \in [n]}, e'$)
if $e' = \perp$ **then**
 if $e = \perp$ **then** $e' \leftarrow 1$
 else $e' \leftarrow e + 1$
else
 req $e' > e \vee e = \perp$
 $G \leftarrow \{uid_1, \dots, uid_n\}$
req $me \neq \perp \wedge me \notin G$
 $AD[\cdot] \leftarrow \perp$
for $i \in [n]$ **do**
 req $*verify-user(uid_i, sig_i)$
 $AD[uid_i] \leftarrow ad_i$
 $uid_{lead} \leftarrow me$
 $e \leftarrow e'$
 $p \leftarrow 0$
 $seed \leftarrow PRG.Init(1^k)$
 $C \leftarrow *encrypt-seed(G, AD)$
 $M \leftarrow (meetingId, e, G, C)$
 $(seed, k) \leftarrow PRG.Eval(seed)$
return M

Algorithm: JoinSession($uid'_{lead}, sig'_{lead}, m, ad$)
req $me \neq \perp \wedge uid'_{lead} \neq me \wedge uid'_{lead} \neq uid_{lead}$
req $*verify-user(uid'_{lead}, sig'_{lead})$
 $uid_{lead} \leftarrow uid'_{lead}$
parse ('epoch', c) $\leftarrow m$
 $(e', p', seed') \leftarrow *decrypt-seed(c, uid'_{lead}, ad)$
if $e \neq \perp$ **then**
 req $e' > e$
 $(e, p) \leftarrow (e', p')$
 $(seed, k) \leftarrow PRG.Eval(seed')$

Helper: $*encrypt-seed(G', AD)$

$C[\cdot] \leftarrow \perp$
for $uid' \in G'$ **do**
 parse (\cdot, \cdot, upk') $\leftarrow uid'$
 $nonce \leftarrow \$_{AEAD.N}$
 if $K[uid'] = \perp$ **then**
 $K[uid'] \leftarrow HKDF((upk')^{usk}, 'KeyMeetingSeed')$
 $ad' \leftarrow (meetingId, me, AD[uid'])$
 $ad'' \leftarrow Hash('EncryptionKeyMeetingSeed' \parallel Hash(ad'))$
 $c' \leftarrow AEAD.Enc(K[uid'], nonce, (e, p, seed), ad'')$
 $C[uid'] \leftarrow (c', nonce)$
return C

Group and key management (leader)

Algorithm: Add($\{(uid_i, ad_i, sig_i)\}_{i \in [n]}, newEpoch$)
req $me \neq \perp \wedge uid_{lead} = me$
 $AD[\cdot] \leftarrow \perp$
for $i \in [n]$ **do**
 req $uid_i \neq me \wedge uid_i \notin G \wedge *verify-user(uid_i, sig_i)$
 $G \leftarrow G \cup \{uid_i\}$
 $AD[uid_i] \leftarrow ad_i$
if newEpoch **then**
 $(e, p) \leftarrow (e + 1, 0)$
 $seed \leftarrow PRG.Init(1^k)$
 $G' \leftarrow G$
else
 $(e, p) \leftarrow (e, p + 1)$
 $G' \leftarrow \{uid_1, \dots, uid_n\}$
 $C \leftarrow *encrypt-seed(G', AD)$
 $(seed, k) \leftarrow PRG.Eval(seed)$
return M $\leftarrow (meetingId, e, G, C)$

Algorithm: Remove($\{uid_i\}_{i \in [n]}$)
req $me \neq \perp \wedge uid_{lead} = me$
for $i \in [n]$ **do** **req** $uid_i \in G$
 $G \leftarrow G \setminus \{uid_1, \dots, uid_n\}$
 $(e, p) \leftarrow (e + 1, 0)$
 $seed \leftarrow PRG.Init(1^k)$
 $AD[\cdot] \leftarrow \perp$
 $C \leftarrow *encrypt-seed(G, AD)$
 $(seed, k) \leftarrow PRG.Eval(seed)$
return M $\leftarrow (meetingId, e, G, C)$

Message processing (participants)

Algorithm: Process(m)
req $me \neq \perp \wedge uid_{lead} \neq \perp \wedge uid_{lead} \neq me$
if $m = ('epoch', c)$ **then**
 $(e', p', seed') \leftarrow *decrypt-seed(c, uid_{lead}, \perp)$
 req $(e', p') = (e + 1, 0)$
 $(e, p) \leftarrow (e', p')$
 $(seed, k) \leftarrow PRG.Eval(seed')$
else if $m = 'period'$ **then**
 $p \leftarrow p + 1$
 $(seed, k) \leftarrow PRG.Eval(seed)$

Helper: $*decrypt-seed(c, uid_{lead}, ad)$

parse (\cdot, id', \cdot, upk') $\leftarrow uid_{lead}$
if $K[uid_{lead}] = \perp$ **then**
 $K[uid_{lead}] \leftarrow HKDF((upk')^{usk}, 'KeyMeetingSeed')$
 $ad' \leftarrow (meetingId, id', ad)$
 $ad'' \leftarrow Hash('EncryptionKeyMeetingSeed' \parallel Hash(ad'))$
 $nonce \leftarrow c$
 $parse(c', nonce) \leftarrow c$
 $parse(e', p', seed') \leftarrow AEAD.Dec(K[uid_{lead}], nonce, c', ad'')$
return $(e', p', seed')$

Helper: $*verify-user(uid', sig')$

parse (meetingId', id', ipk', upk') $\leftarrow uid'$
return meetingId' = meetingId
 $\wedge Sig.Verify(ipk', 'EncryptionKeyAnnouncement', uid', sig')$
 $\wedge PKI.verify-pk(id', ipk')$

Fig. 1: The client protocol of Zoom's cmKEM scheme. The protocol implicitly maintains a state st which includes $me, usk, isk, uid, K, uid_{lead}, G, seed$ as well as the exposed key $st.k$, epoch $st.e$, and period $st.p$.

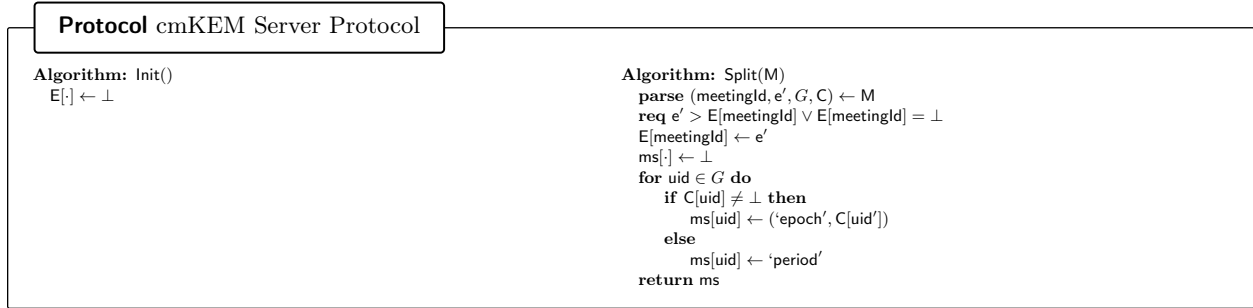


Fig. 2: The server protocol of the cmKEM scheme. The protocol implicitly maintains a state `pub` that just contains the mapping from meetings to their respective latest epoch number received `pub.E[meetingId]`.

uses HKDF for this derivation, which for the purpose of the security analysis we model as a random oracle. For efficiency reasons, this key is cached as part of the sender’s secret state and reused for future messages to or from the same party. The seed is then encrypted using a nonce-based AEAD with a random nonce that is transmitted as part of the resulting ciphertext. The associated data contains the meeting and sender identifiers, as well a fixed context string.

Server protocol. The Split protocol works by delivering the respective AEAD-ciphertext to each party and sending a special “ratchet period” message to parties for which no such ciphertext is specified. In terms of state, the server simply stores the latest epoch for each meeting `meetingId`, which is exposed as `pub.E[meetingId]`. For simplicity, we model that the message $M = (\text{meetingId}, e, G, C)$ sent to the server includes the meeting identifier, the current epoch, as well as set of recipients. Each recipients `uid'` for which C contains a share obtains $m = (\text{'epoch'}, C[\text{uid}'])$, while for other recipients the server delivers $m = \text{'period'}$. The scheme is depicted in Fig. 2.

Security. The following theorem establishes the security of the scheme.

Theorem 1. *Zoom’s cmKEM scheme is secure according to the outlined definition under the Gap-DH assumption, if the AEAD scheme is secure, Hash collision resistant, the signature scheme is EUF-CMA secure, the PRG satisfies the standard indistinguishability from random notion, and HKDF is modeled as a random oracle.*

A full proof is presented in Appendix C.4. In short, based on the security of Gap Diffie-Hellman, we can first switch to a hybrid where we use independently generated symmetric keys as opposed to the outputs of the DH operation (between the leader and each participant), programming the random oracle to make things look consistent on corruption. Then, we can argue that each of the adversary’s winning conditions in the game cannot be triggered based on the unforgeability of the signature scheme (credentials cannot be forged), the authenticity of the AEAD (malicious keys cannot be injected), and the confidentiality of the AEAD (encrypted keys cannot be distinguished from encryptions of random messages).

According to the whitepaper [11], Zoom’s scheme performs Diffie-Hellman over Curve25519.⁷ We note that the Gap-DH assumption (rather than, e.g., CDH) appears to be rather intrinsic to this kind of simple Diffie-Hellman based protocol and has been assumed for Curve25519 before [10,1]. Moreover, Zoom uses XChaCha20Poly1305 with 192-bit nonces as the nonce-based AEAD scheme, and HKDF for both the key-derivation as well as the PRG. (We model the latter use as a PRG to clarify the exact required security properties.) Finally, for a signature scheme, Zoom uses EdDSA over Ed25519 satisfying EUF-CMA security [13].

⁷ Technically, Curve25519 breaks the abstraction of cyclic groups we use (for simplicity) to present our scheme. We refer to the analysis of the HPKE standard [1] for an extended discussion and the formalization of *nominal groups* with the respective Gap-DH assumption. Their results directly apply to our construction.

3 Leader-based GCKA with Liveness

We now abstract the core of Zoom’s E2EE meetings protocol¹ as a *leader-based continuous group key agreement with liveness* (LL-CGKA) scheme. On a high level, the primitive works similarly to the previously introduced cmKEM one, with the following differences: (1) participants are aware of the group roster and in particular only use keys for which they know the roster, (2) as a result participants can no longer run ahead of their leader in terms of the period, and (3) liveness is enforced.

Liveness. To achieve liveness, the LL-CGKA primitive is time based. More concretely (1) algorithms can depend on time and (2) in addition to event-based actions (e.g., reacting to an incoming packet), there are also time-driven actions. We make the following (simplifying) assumptions:

- Each party has a *local clock*, and all clocks run at *the same speed* (constant drift).
- Local algorithms complete instantaneously, i.e., no time elapses between invocation and completion. As a consequence, the algorithms simply take the party’s current time as an input argument.

We remark that the vast majority of Zoom meetings last only a couple of hours, which limits any significant amount of clock drift in practice and thus justifies the first assumption.

3.1 Syntax

The algorithms of a LL-CGKA scheme closely follow the ones of a cmKEM scheme, with two major differences. First, client algorithms take the current *local time* as input. Second, there are *clock ticking* algorithms that allow to specify clock-driven actions, i.e., actions that happen at a certain time rather upon receiving a message.

As in cmKEM, the server performs message routing. Additionally, it hands out the current public⁸ group state to newly joining parties, freeing the leader from maintaining additional state.

Definition 2. A LL-CGKA scheme consists of the algorithms described in the following, where, for ease of presentation, the client state ust is assumed to expose the following fields:

- The user’s current epoch ust.e and period ust.p .
- For each epoch and period a key $\text{ust.k}[e, p]$ (or \perp if not known yet). It is assumed that operations do not change keys once they are defined.
- The user’s current view on the group ust.G .

User management:

- $(\text{ust}, \text{uid}, \text{sig}) \leftarrow \text{CreateUser}(\text{time}, \text{id}, \text{meetingId})$ creates an ephemeral user for the given identity and meeting. Outputs the user’s initial state ust , their identity uid , and credential sig binding uid to id .
- $\text{id} \leftarrow \text{Identity}(\text{uid})$ and $\text{meetingId} \leftarrow \text{Meeting}(\text{uid})$ are deterministic algorithms that return the ephemeral user’s long-term identity and meeting, respectively.
- $\text{ust}' \leftarrow \text{CatchUp}(\text{ust}, \text{time}, \text{grpPub})$ prepares the user for joining the group by processing the current public group state grpPub provided by the sever.

Leader’s algorithms:

- $(\text{ust}', M) \leftarrow \text{Lead}(\text{ust}, \text{time}, \{(\text{uid}_i, \text{sig}_i)\}_{i \in [n]})$ instructs the user to become the new group leader with the specified participants. Outputs a message to be split and distributed to the other group members.
- $(\text{ust}', M) \leftarrow \text{Add}(\text{ust}, \text{time}, \{(\text{uid}_i, \text{sig}_i)\}_{i \in [n]})$ is used to add users uid_1 to uid_n to the group.
- $(\text{ust}', M) \leftarrow \text{Remove}(\text{ust}, \text{time}, \{\text{uid}_i\}_{i \in [n]})$ is used to remove users uid_1 to uid_n from the group.
- $(\text{ust}', M) \leftarrow \text{LeaderTick}(\text{ust}, \text{time})$ is executed on each clock tick by the leader. Outputs the leader’s updated state and an optional messages M .

⁸ By public, we mean known to the (untrusted) Zoom server; i.e., the current roster, but not any keys.

Participants' algorithms:

- $ust' \leftarrow \text{Follow}(ust, \text{time}, m, \text{uid}'_{\text{lead}}, \text{sig}'_{\text{lead}})$ instructs the user to treat $\text{uid}'_{\text{lead}}$ as the new leader. Expects the first message share m from the new leader.
- $ust' \leftarrow \text{Process}(ust, \text{time}, m)$ is used by participants to process any incoming message m .
- $(ust', \text{alive}, \text{sig}') \leftarrow \text{ParticipantTick}(ust, \text{time})$ is executed by a participant on each clock tick. The flag alive indicates whether the participant is still in the meeting or dropped out (for a violation of liveness) and optionally updates the credential (for the server) with $\text{sig}' = \perp$ denoting no update.

Server's algorithms:

- $\text{pub} \leftarrow \text{Init}()$ generates an initial server state.
- $(\text{pub}', \{(uid_i, m_i)\}_{i \in [n]}) \leftarrow \text{Split}(\text{pub}, M)$ is a deterministic algorithm that takes message M and splits out each user uid 's share m .
- $\text{grpPub} \leftarrow \text{GroupState}(\text{pub}, \text{meetingId})$ is a deterministic algorithm that returns the public group state.

We discuss correctness, and in particular how it is affected by the liveness properties, in Appendix D.4.

Meeting Flow. Let us briefly discuss how Zoom uses the above defined LL-CGKA abstraction to orchestrate a meeting. Parties first generate a per-meeting ephemeral identity using `CreateUser` and communicate it to the Zoom server. Before any party can start or join the meeting, the server hands them the most recent public group state (using `GroupState`) that the party processes using `CatchUp`.

To start a meeting, Zoom then instructs the initial host to invoke the `Lead` algorithm. Later, the server can instruct the leader to modify the set of meeting participants using `Add` and `Remove`. All messages sent by the leader to current participants are mediated through the server, which uses `Split` to compute the share of the message that each participant needs.

Participants join the meeting by invoking `Follow` with their respective share of the message generated by the leader's corresponding `Lead` (or `Add`) invocation. They also use `Process` when receiving further messages from the same leader. Observe that, in some cases, the `Follow` invocation might not be enough for participants to know about the group roster and the latest key. Instead, it might take up to an additional message generated by `LeaderTick` for the participant to fully join the meeting.

Analogously, to switch leaders, the new one is instructed to invoke `Lead` and all other participants are instructed to invoke `Follow` again. Note that it is not required for the new leader to already be part of the group — the `CatchUp` algorithm can directly be followed by `Lead` (instead of `Follow`) to immediately start as the new leader.

3.2 Security Definition

Overall, the game follows closely the one of the `cmKEM` primitive outlined in Section 2.3. In the following we discuss the key aspects and highlight the differences to the `cmKEM` game. We refer to Appendix D.1 for a formal definition.

Clocks. The security game maintains a global clock time. Each honestly created user uid maintains a local clock that is specified as an offset to the global one; that is, all local clocks run at the same speed. (For our analysis, we do not make use of the fact that two users uid and uid' belonging to the same long-term identity id , i.e. a device, typically would have the same local clock. Trivially, our results also hold for this special case.)

The adversary chooses each user's offset and drives the global clock, i.e., decides whenever the clock is supposed to advance by a tick. Those ticks model an abstract discrete unit of time, which can be thought of as milliseconds or nanoseconds, roughly corresponding to the precision of clocks used by the various parties. Whenever the adversary ticks the global clock, each party's local clock thus also advances, and their respective procedures `LeaderTick` or `ParticipantTick` are invoked, depending on whether the party is currently a leader or not.

Liveness. An important objective of the LL-CGKA primitive is to ensure liveness: all participants must either keep up with the current meeting’s state or drop out of the meeting. This is formalized as follows: whenever `ParticipantTick` indicates that `uid` is still alive, then the participant’s state must not be too outdated, which in turn is defined as that the participant’s current leader *must have been in the same state recently*. How recently, exactly, is a parameter of our security definition we call the *liveness slack*; we introduce the concrete slack achieved by Zoom’s protocol as part of its description below.

Observe that this formalization essentially means that whenever the leader makes a change to the group by either adding or removing parties (resulting in an epoch or period change), this change cannot be withheld by a malicious server. We thus call this property *key liveness* and briefly discuss *content liveness* in Section 5.

Confidentiality. Group key confidentiality is formalized analogously as in the cmKEM scheme. That is, upon a challenge, the security game outputs either the real or an independent uniform random key depending on a bit b . We remark that the game only allows to challenge keys for epochs and periods that are to be used in the higher-level application, omitting those that are skipped by the LPL mechanism and hence never output. This simplifies the notion as, in contrast to the cmKEM security notion, each challengeable key has a well-defined group roster associated.

Consistency, authenticity, and no-merging. The game ensures both *key consistency* and *group consistency*, meaning that for a given honest (and uncompromised) leader, epoch, and period, all (uncompromised) participants agree on a key and group roster. However, a malicious server could cause the group to split by assigning different leaders to different partitions of the group, in which case those partitions will no longer agree on the key. Furthermore, two parties, say Alice and Bob, in different partitions might both believe to have a third party Charlie in their group, or even believe to be in the same group with the other party – that is, group rosters output by different partitions are not guaranteed to be disjoint. However, the game ensures that after such an aforementioned (inherent) splitting attack, the various partitions cannot be re-merged into a consistent state, which makes the attack easier to detect.

Remark 1 (Insider security). This work focuses on outsider security, and only formalizes limited insider security guarantees. For instance, whenever the adversary performs a trivial injection, enabled by, e.g. a corruption of the leader, most security properties are (temporarily) disabled. On the other hand, we do formalize that confidentiality and authenticity recover after switching from a malicious leader to an honest one. While Zoom’s protocol does not aim to provide strong guarantees in the presence of malicious insiders (as full insider security would, e.g. require asymmetric authentication for video data), a more comprehensive analysis of the properties it does achieve would nevertheless be interesting.

3.3 Zoom’s Scheme

We now describe Zoom’s LL-CGKA scheme. On a high-level, the protocol enhances the cmKEM scheme by having the leader broadcast the group membership to all participants, as well as regular *heartbeat messages* that help guarantee liveness. A formal description is presented in Fig. 3. Additional details can be found in Appendix D.2.

Leader Participant List. For the participants to learn the group roster, the session leader broadcasts the so-called *leader participant list (LPL)* tabulating the members. The LPL is, for bandwidth efficiency, represented as a linked list of differential updates containing the set of added and removed participants since the last LPL. Each message also references the leader’s current epoch e and period p . For efficiency reasons, an LPL message is not sent on every single change to the group roster, but on regular intervals instead. (It is skipped if no change to the group roster has been done in the meantime.) To ensure that parties know to whom they speak to, the scheme only proceeds to epochs and periods that have been certified by an LPL message. Re-keying is nevertheless performed eagerly, potentially leading to unused keys.

The protocol furthermore relies on the LPL to communicate the group roster to newly joining parties. To avoid having new parties process the entire history of LPL messages, thus increasing the server’s storage

Protocol Zoom's Client LL-CGKA

User management

Algorithm: CreateUser(time, id, meetingId)
 $(st, me, sig) \leftarrow \text{cmKEM.CreateUser}(id, \text{meetingId})$
 $(isk, \cdot) \leftarrow \text{PKI.get-sk}(id)$
 $\text{lastHb} \leftarrow \text{time}$
 $\text{uid}_{\text{lead}} \leftarrow \perp$
 $G \leftarrow \emptyset$
 $e, p, e_{\text{next}}, p_{\text{next}}, v, t \leftarrow 0$
 $e_{\text{pub}} \leftarrow \perp$
 $k[\cdot, \cdot] \leftarrow \perp$
 $\delta[\cdot] \leftarrow \infty$
 $\text{lplHash}, \text{hbHash} \leftarrow \perp$
return (me, sig)

Algorithm: Identity(uid)
return $\text{cmKEM.Identity}(uid)$

Algorithm: Meeting(uid)
return $\text{cmKEM.Meeting}(uid)$

Algorithm: CatchUp(time, grpPub)
req $\text{uid}_{\text{lead}} = \perp$
parse $(e_{\text{pub}}, \text{lpls}', \text{hb}') \leftarrow \text{grpPub}$
// Process LPLs
while $\text{lpls}' \neq \{\}$ **do**
 $\text{lpl}' \leftarrow \text{lpls}'.\text{deq}()$
 try **receive-LPL*(lpl')
// Store Heartbeat (no verification)
 $\text{hbHash} \leftarrow \text{Hash}(\text{hb}')$

Participants' algorithms

Algorithm: Follow(time, m' , $\text{uid}'_{\text{lead}}$, $\text{sig}'_{\text{lead}}$)
req $\text{uid}'_{\text{lead}} \neq \perp \wedge \text{uid}'_{\text{lead}} \neq me$
parse $(m'_K, \text{lpl}', \text{hb}') \leftarrow m'$
if $\text{uid}'_{\text{lead}} \neq \perp$ **then**
 req $\text{lpl}' \neq \perp \wedge \text{hb}' \neq \perp$
try $st \leftarrow \text{cmKEM.JoinSession}(st, \text{uid}'_{\text{lead}}, \text{sig}'_{\text{lead}}, m'_K, \perp)$
 $\text{Key}[st.e, st.p] \leftarrow st.k$
 $\text{uid}_{\text{lead}} \leftarrow \text{uid}'_{\text{lead}}$
 $e_{\text{pub}} \leftarrow \perp$
if $\text{lpl}' \neq \perp$ **then**
 try **receive-LPL*(lpl')
 req $me \in G \wedge \text{hb}' \neq \perp \wedge (e_{\text{next}}, p_{\text{next}}) = (st.e, st.p)$
if $\text{hb}' \neq \perp$ **then**
 try **receive-heartbeat*(hb')

Algorithm: Process(time, m')
req $\text{uid}_{\text{lead}} \neq \perp \wedge \text{uid}_{\text{lead}} \neq me$
parse $(m'_K, \text{lpl}', \text{hb}') \leftarrow m'$
if $m'_K \neq \perp$ **then**
 try $st \leftarrow \text{cmKEM.Process}(st, m'_K)$
 $\text{Key}[st.e, st.p] \leftarrow st.k$
if $\text{lpl}' \neq \perp$ **then**
 try **receive-LPL*(lpl')
 req $me \in G \wedge \text{hb}' \neq \perp$
if $\text{hb}' \neq \perp$ **then**
 try **receive-heartbeat*(hb')

Leader's algorithms

Algorithm: Lead(time, $\{(uid_i, sig_i)\}_{i \in [n]}$)
if $e_{\text{pub}} \neq \perp$ **then** $e' \leftarrow e_{\text{pub}} + 1$
else $e' \leftarrow \perp$
try $(st, M_K) \leftarrow \text{cmKEM.StartSession}(st, \{(uid_i, \perp, sig_i)\}_{i \in [n]}, e')$
 $\text{uid}_{\text{lead}} \leftarrow me$
 $(e, p) \leftarrow (st.e, st.p)$
 $e_{\text{pub}} \leftarrow \perp$
 $G \leftarrow \{uid_1, \dots, uid_n\} \cup \{me\}$
 $\text{Added}, \text{Removed} \leftarrow \emptyset$
 $\text{numLplLinks} \leftarrow \text{max-links}$
 $\text{lpl} \leftarrow \text{*send-LPL}()$
 $\text{hb} \leftarrow \text{*send-heartbeat}()$
 $M \leftarrow (me, M_K, \text{lpl}, \text{hb})$
return M

Algorithm: Add(time, $\{(uid_i, sig_i)\}_{i \in [n]}$)
req $\text{uid}_{\text{lead}} = me$
for $i \in [n]$ **do**
 req $uid_i \notin G$
 $G \leftarrow G \cup \{uid_i\}$
 $\text{Added} \leftarrow \text{Added} \cup \{uid_i\}$
if $p \geq p_{\text{MAX}}$ **then**
 try $(st, M_K) \leftarrow \text{cmKEM.Add}(st, \{(uid_i, \perp, sig_i)\}_{i \in [n]}, \text{true})$
 $(e, p) \leftarrow (st.e, st.p)$
else
 try $(st, M_K) \leftarrow \text{cmKEM.Add}(st, \{(uid_i, \perp, sig_i)\}_{i \in [n]}, \text{false})$
 $M \leftarrow (me, M_K, \perp, \perp)$
return M

Algorithm: Remove(time, $\{uid_i\}_{i \in [n]}$)
req $\text{uid}_{\text{lead}} = me$
for $i \in [n]$ **do**
 req $uid_i \in G \wedge uid_i \neq me$
 $G \leftarrow G \setminus \{uid_1, \dots, uid_n\}$
 $\text{Removed} \leftarrow \text{Removed} \cup \{uid_1, \dots, uid_n\}$
 $\text{Added} \leftarrow \text{Added} \setminus \{uid_1, \dots, uid_n\}$
 try $(st, M_K) \leftarrow \text{cmKEM.Remove}(st, \{uid_i\}_{i \in [n]})$
 $(e, p) \leftarrow (st.e, st.p)$
 $M \leftarrow (me, M_K, \perp, \perp)$
return M

Time driven

Algorithm: LeaderTick(time)
req $\text{uid}_{\text{lead}} = me$
 $\text{lpl}, \text{hb} \leftarrow \perp$
if $\text{time} - \text{lastHb} \geq \Delta_{\text{LPL}}$ **then**
 if $\text{Added} \neq \emptyset \vee \text{Removed} \neq \emptyset$ **then**
 $\text{lpl} \leftarrow \text{*send-LPL}()$
 $\text{hb} \leftarrow \text{*send-heartbeat}()$
 else if $\text{time} - \text{lastHb} \geq \Delta_{\text{heartbeat}}$ **then**
 $\text{hb} \leftarrow \text{*send-heartbeat}()$
 $M \leftarrow (me, \perp, \text{lpl}, \text{hb})$
return M

Algorithm: ParticipantTick(time)
req $\text{uid}_{\text{lead}} \neq \perp \wedge \text{uid}_{\text{lead}} \neq me$
 $\text{alive} \leftarrow (\text{time} - \text{lastHb} \leq \Delta_{\text{live}})$
return (alive, \perp) *// no updated credential*

Fig. 3: The client part of Zoom's overall LL-CGKA scheme. The description omits the state ust , which is input and output by most algorithms as specified in Section 3.1, and implicitly updated. ust is described in detail in Appendix D.2.

Protocol Zoom's Client LL-CGKA Helpers

```

Helper: *send-LPL()
v ← v + 1
if numLplLinks ≥ max-links then
  lpl ← (me, v, true, ⊥, G, ⊥, e, p)
  numLplLinks ← 1
else
  lpl ← (me, v, false, lplHash, Added, Removed, e, p)
  numLplLinks ← numLplLinks + 1
lplHash ← Hash(lpl)
(Added, Removed) ← ∅
return lpl

Helper: *receive-LPL(lpl')
parse (uidlead', v', coalesced', lplHash',
      Added', Removed', e', p') ← lpl'
req uidlead' = uidlead
req v = ⊥ ∨ v' = v + 1
v ← v'
if ¬coalesced then
  req lplHash' = lplHash ∧ lplHash ≠ ⊥
  req (Removed' \ G) = ∅
  G ← G \ Removed'
  req (Added' ∩ G) = ∅
  G ← G ∪ Added'
else
  G ← Added'
lplHash ← Hash(lpl')
(enext, pnext) ← (e', p')

Helper: *send-heartbeat()
t ← t + 1
sighb ← Sig.Sign(isk, 'LeaderParticipantList',
                  (me, t, hbHash, time, v, lplHash, e, p))
hb ← (t, time, sighb)
hbHash ← Hash(hb), lastHb ← time
return hb

Helper: *receive-heartbeat(hb)
parse (t', time', sighb') ← hb
req t = ⊥ ∨ t' = t + 1
t ← t'
(·, ipk') ← cmKEM.Identity(uidlead)
req Sig.Verify(ipk', 'LeaderParticipantList',
              (uidlead, t, hbHash, time', v, lplHash, enext, pnext), sighb')
req Key[enext, pnext] ≠ ⊥
(e, p) ← (enext, pnext)
*update-drift(time')
*update-liveness(time')
hbHash ← Hash(hb)
t ← t'

Helper: *update-drift(time')
δ[uidlead] ← min(δ[uidlead], time - time')

Helper: *update-liveness(time')
lastHb ← time' + δ[uidlead]

```

Fig. 4: Helper algorithms for the client part of Zoom's overall LL-CGKA scheme.

requirement, the leader will from time to time use a special *coalesced* LPL message encoding the entire group.⁹ A joining party therefore needs all the links up to and including the latest coalesced message only. The frequency of coalesced messages is determined by the parameter *max-links*.

Heartbeats. The LPL messages are unauthenticated. To authenticate them, the leader broadcasts a signature (of a hash) thereof under the leader's long-term identity key. Those signatures moreover form another hash chain, with each signature including the hash of the previous one to ensure the continuity of the meeting. That is, while certain splitting attacks — where a malicious server might tell subgroups to accept different leaders — are unavoidable, those diverging meetings cannot later be merged.

The leader broadcasts one such signature at least at a fixed interval $\Delta_{\text{heartbeat}}$, even when no LPL has been sent for lack of any change to the group membership. Since they are regularly sent, these signatures are called *heartbeat messages* and double as a mechanism to ensure liveness. To this end, the signature additionally includes the latest epoch *e*, and period *p*. Hence, if an attacker attempts to withhold either key rotations or updates to the membership, causing a participant to be stuck in an old state, they would need to withhold the heartbeat message as well. As a countermeasure, participants drop out from the meeting if they have not received a heartbeat message for a certain amount of time.

For this mechanism to not abruptly end meetings (despite potential network hiccups), participants do not expect to receive the heartbeats in perfectly regular intervals. Rather, each heartbeat itself contains a timestamp *time'* (the sending time) whose state it certifies. Receiving this heartbeat then prolongs the liveness of the receiver until time $\text{time}' + \delta + \Delta_{\text{live}}$, when the party will drop out if no further heartbeat has been received. Here, δ denotes an upper bound on the clock drift (between the participant and their respective leader) and Δ_{live} denotes a protocol parameter. As a best effort to prevent this from happening, the server will elect a new leader whenever the current one struggles to upload heartbeats.

⁹ In the deployed version of the protocol, the coalesced LPL also includes a list of all participants who were in the meeting at some point in the past but have since left. This additional information is displayed in the client's user interface, but is not modeled in this work.

The protocol estimates the upper bound on the clock drift δ as follows: Upon receiving the first heartbeat with timestamp t' at local time t from a given leader, the protocol simply assumes that $t - t'$ is the drift, i.e., that the heartbeat has been delivered instantaneously. Clearly, $t' + \delta = t$ is an upper bound on the effective sending time. Upon receiving a subsequent heartbeat, the party corrects the drift to $t - t'$ whenever this is smaller, and otherwise keeps it unchanged. Hence, if the network delay, and thus the interval between received heartbeat, increases (e.g., due to a network attacker) then each subsequently received heartbeat extends liveness by a smaller amount, until the party eventually drops out. Conversely, if the network delay decreases, the drift estimates and, hence, the liveness assurances improve.

Evolving the group. The protocol uses the cmKEM scheme to rotate keys whenever the group membership changes. The participants, however, do not immediately transition to the new epoch or period upon receiving such a cmKEM message. Rather, they just store the new key. Only once the group membership is known via receiving a corresponding LPL message and heartbeat, they transition to the new epoch and period and advertise the respective key for content encryption to the higher-level protocol. For membership changes containing only additions, the protocol avoids overly frequent epoch changes by rotating the period instead, however, limiting the number of consecutive periods to a fixed number p_{MAX} .

Joining a meeting. To join a meeting, a party needs to learn the latest key and group roster in an authenticated manner. The former is communicated via a cmKEM message and the latter via the sequence of LPL messages starting with the latest coalesced one. Authentication of the LPL is achieved by verifying the latest heartbeat message that certifies the final LPL message, as well as epoch and period numbers. (The previous links are implicitly authenticated due to the links forming a hash chain.)

Leader changes. A newly elected leader continues the meeting by starting a new cmKEM session and generating a coalesced LPL message and a heartbeat, which the server then distributes to the other participants. The new leader will continue the relevant counters (i.e., e , p , and t) and hash chains where the old leader left off, such that they uniquely identify a meeting state. The server is responsible to ensure that the party has the latest state the moment it becomes the new leader.

Note that the new leader obtains the group roster from the server rather than deducing it from the previous LPL messages. Otherwise, they might inadvertently revert some of the previous leader's final changes to the group, if for instance the previous leader added or removed a party on the cmKEM level but did not manage to broadcast a corresponding LPL message before dropping out. Users are shown a warning on every leader change, and are advised to manually check whether the group roster displayed in their client is expected.

The server scheme. The messages the leader uploads consist of up to three components, a cmKEM message, a LPL and a heartbeat message. If the message contains a cmKEM message, then the server splits this using the respective cmKEM algorithm and forwards the respective share alongside the LPL and heartbeat (if present) to the users. Otherwise, the server forwards the LPL and heartbeat messages to the last known roster, as derived from the cmKEM messages. See Appendix D.2 for details.

Security. Security is summarized in our main result below, with a more detailed proof given in Appendix D.3.

Theorem 2. *Zoom's LL-CGKA scheme is secure with the liveness slack of P being at most*

$$\min(n \cdot \Delta_{\text{live}}, t_{\text{now}} - t_{\text{joined}}) + \Delta_{\text{live}},$$

where t_{now} denotes the current time, t_{joined} the time P joined the meeting, and n denotes the number of distinct leaders P has encountered so far. Liveness holds if all those leaders have followed the protocol, while all other properties hold as long as the current leader is honest.

Proof (Sketch). Confidentiality and key consistency follow directly from the underlying cmKEM scheme which is used to distribute the group keys. While the LL-CGKA notion mandates slightly stronger properties, those additional assurances relate directly to members only transitioning to subsequent periods if their

leader initiated this. This is ensured by parties only transitioning to a new state once a heartbeat certified it, leveraging the unforgeability of the employed signature scheme. Similarly, group consistency — i.e., authenticity of each participant’s view on the group roster — is ensured by the combined LPL and heartbeat mechanism, with the LPL distributing the group and the heartbeat authenticating the LPL. Additionally, the hash links of the heartbeat messages yields the no-merging property after a group-splitting attack.

Finally, observe that liveness slack is directly linked to the accuracy of each party’s estimate on the clock drift with their respective leader: If the estimate were precise, then each party would have a liveness slack of at most Δ_{live} since they would know exactly when the last heartbeat they received has been sent allowing them to drop out Δ_{live} after. Further, the estimate only degrades by at most Δ_{live} with each leader change — the maximum interval between receiving the old leader’s last heartbeat and the new leader’s first one. \square

The above theorem relies on the underlying cmKEM scheme being secure according to the respective definition, the signature scheme being EUF-CMA secure, and the hash function being collision resistant. According to the whitepaper [11], Zoom’s instantiation uses SHA256 and EdDSA (as provided by `libsodium`) for the hash function and signature algorithm, respectively, satisfying those requirements [19,13].

Concrete parameters. At the time of writing, Zoom uses $\Delta_{\text{live}} = 100s$, $\Delta_{\text{heartbeat}} = 10s$, $\Delta_{\text{LPL}} = 2s$, and $\text{max-links} = 20$, respectively. Moreover, $\text{p}_{\text{MAX}} = 0$, i.e., Zoom always ratchets the full epoch instead of the period³.

4 Improved Liveness

4.1 Limitations of Zoom’s Protocol

For a typical meeting with a single (honest) host that stays online for the duration of the entire meeting — and thus is the leader for the entire meeting — Zoom’s current scheme¹ provides strong liveness properties. Indeed, to the best of our knowledge, Zoom is the only E2EE group video protocol that provides any such liveness assurance. As highlighted by Theorem 2, however, there two distinct aspects with respect to which the assurances could be further improved:

1. Zoom’s current liveness assurance degrade in the number of meeting leaders encountered. This is sub-optimal for a protocol such as Zoom where the (untrusted) server can initiate leader changes.¹⁰
2. While all other security properties, such as key confidentiality and authenticity, recover after removing a malicious party from the meeting, liveness does not.¹¹

We particularly stress that both aspects are not merely deficiencies of our analysis. Concrete (though contrived) attacks exist, even if they could be mitigated by countermeasures relying on the end user, such as user-interface warnings.

Lemma 1. *Even with all honest participants, the liveness properties of Zoom’s LL-CGKA scheme degrade in the number of leader changes, assuming an all powerful malicious server carefully orchestrating the meeting.*

Proof. Consider a meeting with parties P_1, P_2, \dots, P_n , as well as a designated party P^* . All parties, unbeknownst to each other, have precisely synchronized clocks. The party P_1 is the one to start the meeting and act as its initial leader. When adding the parties P_2, \dots, P_n to the meeting, the network adversary delivers the respective messages immediately. That is, the moment those parties create their ephemeral user identities $\text{uid}_2, \dots, \text{uid}_n$, party P_1 is immediately instructed to add them to the meeting using `Add` producing `M`, and the respective shares obtained by `Split` are handed to the parties to execute `Follow` without any delay. (To

¹⁰ This is currently remedied by the client showing a warning upon each leader change, since the leader-authentication codes anyway require to repeat the authentication process in this event. With the introduction of the advanced PKI replacing the leader-authentication codes, Zoom might however consider dropping those warnings.

¹¹ Note that Zoom does not aim to provide strong guarantees *while* a malicious insider is part of the meeting. Yet, removing a malicious party should ideally reestablish security without the need to restart the entire meeting.

this end, assume that the heartbeat interval perfectly aligns with the moment all those parties join.) This results in each of those parties estimating their drift to be 0, i.e., $\delta_{\text{uid}_j}[\text{uid}_1] = 0$ for $j \in \{2, \dots, n\}$.

In contrast, when party P^* joins the meeting, their respective ephemeral identity uid^* is still handed immediately to P_1 , but the respective response delayed by Δ_{live} . Assuming P^* created their identity at time t and got the LL-CGKA message at time $t + \Delta_{\text{live}}$, but with timestamp t , then P^* assumes that their clock runs ahead by Δ_{live} , i.e., $\delta_{\text{uid}^*}[\text{uid}_1] = \Delta_{\text{live}}$. All subsequent heartbeats from P_1 are then delivered to P^* with a delay of Δ_{live} . As a result, if P_1 sends a further heartbeat at time t' , P^* will set $\text{lastHb} \leftarrow t' + \Delta_{\text{live}}$ and therefore extend the time until they drop out until $t' + 2\Delta_{\text{live}}$ (instead of the optimal $t' + \Delta_{\text{live}}$).

Next, consider P_1 sending their last heartbeat (which is delivered to all parties as previously described) at a time t_2 after P^* joined the meeting, and immediately afterwards the party P_2 becoming the leader (still at time t_2). Again, the messages derived from the output of `Lead` are distributed to P_3, \dots, P_n without delay, again resulting in $\delta_{\text{uid}_j}[\text{uid}_2] = 0$. For P^* , on the other hand, P_1 's last heartbeat is delivered at time $t_2 + \Delta_{\text{live}}$, extending liveness until $t_2 + 2\Delta_{\text{live}}$. The adversary now takes advantage of this fact by delaying the first message from P_2 as well as all subsequent ones by $2\Delta_{\text{live}}$. This process can then be repeated with sequentially switching leaders to P_3, P_4, \dots, P_n , leading to a liveness slack of $(n + 1)\Delta_{\text{live}}$. \square

Lemma 2. *If parties join a meeting that currently has a malicious leader colluding with a party with extensive control over Zoom's server infrastructure, then the liveness assurance can be arbitrarily broken even after all malicious parties have been removed from the meeting (and an honest leader has taken over).*

Proof. Consider a malicious insider attacker P_M starting a meeting. Moreover, assume that there are two honest parties P_A and P_B , where first P_A wants to join and at a later point P_B wants to join. Assume that all have perfectly synchronized clocks. In the meeting, attacker first adds P_A to the group, without any delay, i.e., such that $\delta_{\text{uid}_A}[\text{uid}_M] = 0$. At time t , right when P_B is about to join (e.g., once P_B advertised their ephemeral uid_B) the malicious insider does the following:

1. P_M creates k heartbeat messages $t + 1, t + 2, \dots, t + k$ (when t denotes the number of heartbeats created so far) for which they pretend to be normally spaced out by $\Delta_{\text{heartbeat}}$ with respect to the included timestamps.
2. P_M then adds P_B to the meeting in state $t + k$, i.e., the first heartbeat signing over the LPL containing uid_B is with counter $t + k + 1$.

The attacker controlling Zoom's server infrastructure now delivers those messages as follows:

1. Immediately deliver the welcoming message, including the $(t + k + 1)$ -th heartbeat, at time t to P_B . As a result P_B will set $\delta_{\text{uid}_B}[\text{uid}_M] = -k \cdot \Delta_{\text{heartbeat}}$, since to P_B it looks like the clock of P_M simply runs ahead.
2. Immediately make P_B the new leader at time t .
3. Deliver all the k intermediate heartbeats to P_A at the regular interval $\Delta_{\text{heartbeat}}$. At time $t + k \cdot \Delta_{\text{heartbeat}}$ first deliver the messages corresponding to P_B joining and then, immediately afterwards, the first message from the new leader P_B .

It is easy to see that P_A does not drop out as they get heartbeats exactly as if the meeting would progress normally. More concretely, to P_A it looks like a perfectly normal meeting in which P_B joins at time $t + k \cdot \Delta_{\text{heartbeat}}$. At the end, P_A will still accept the message from P_B , thinking that the clock of P_B must run ahead. \square

As a result, we now propose two alternative strengthened liveness protocols.

4.2 Additional Interaction

As a first proposal we suggest adding additional interaction in the form of sporadic messages from each participant.

Protocol Client LL-CGKA (Improvement 1)

User management

Algorithm: CreateUser(time, id, meetingId)
 $(st, me, sig') \leftarrow \text{cmKEM.CreateUser}(id, \text{meetingId})$

```

nonce  $\leftarrow \mathcal{N}$ 
nonce'  $\leftarrow \perp$ 
lastNonce  $\leftarrow \text{time}$ 
sig  $\leftarrow (\text{sig}', \text{nonce})$ 

```

$(\text{isk}, \cdot) \leftarrow \text{PKI.get-sk}(id)$

lastHb $\leftarrow \text{time}$

uid_{lead} $\leftarrow \perp$

$G \leftarrow \emptyset$

$e, p, e_{\text{next}}, p_{\text{next}}, v, t \leftarrow 0$

$e_{\text{pub}} \leftarrow \perp$

$k[\cdot, \cdot] \leftarrow \perp$

$\delta[\cdot] \leftarrow \infty$

lplHash, hbHash $\leftarrow \perp$

return (me, sig)

Participants' algorithms

Algorithm: Follow(time, m', uid'_{lead}, sig'_{lead})

req uid_{lead} $\neq \perp \wedge \text{uid}'_{\text{lead}} \neq \text{me}$

parse (m'_K, lpl', hb') $\leftarrow m'$

if uid_{lead} $\neq \perp$ **then**

req lpl' $\neq \perp \wedge \text{hb}' \neq \perp$

st' $\leftarrow \text{cmKEM.JoinSession}(st, \text{uid}'_{\text{lead}}, \text{sig}'_{\text{lead}}, m'_K, \text{nonce})$

if st' = \perp **then** // joining failed

// try previous nonce

try st $\leftarrow \text{cmKEM.JoinSession}(st, \text{uid}'_{\text{lead}}, \text{sig}'_{\text{lead}}, m'_K, \text{nonce}')$

else

st $\leftarrow \text{st}'$

Key[st.e, st.p] $\leftarrow \text{st.k}$

uid_{lead} $\leftarrow \text{uid}'_{\text{lead}}$

$e_{\text{pub}} \leftarrow \perp$

if lpl' $\neq \perp$ **then**

try *receive-LPL(lpl')

req $me \in G \wedge \text{hb}' \neq \perp \wedge (e_{\text{next}}, p_{\text{next}}) = (st.e, st.p)$

if hb' $\neq \perp$ **then**

try *receive-heartbeat(hb')

Algorithm: ParticipantTick(time)

req uid_{lead} $\neq \perp \wedge \text{uid}'_{\text{lead}} \neq \text{me}$

alive $\leftarrow (\text{time} - \text{lastHb} \leq \Delta_{\text{live}})$

if time - lastNonce $\geq \Delta_{\text{nonce}}$ **then**

nonce' $\leftarrow \text{nonce}$

nonce $\leftarrow \mathcal{N}$

lastNonce $\leftarrow \text{time}$

sig $\leftarrow (\text{sig}', \text{nonce})$

return (alive, sig)

else

return (alive, \perp)

Leader's algorithms

Algorithm: Lead(time, {(uid_i, sig_i)}_{i ∈ [n]})

if $e_{\text{pub}} \neq \perp$ **then** $e' \leftarrow e_{\text{pub}} + 1$

else $e' \leftarrow \perp$

for $i \in [n]$ **do**

parse (sig'_i, nonce_i) $\leftarrow \text{sig}_i$

try (st, M_K) $\leftarrow \text{cmKEM.StartSession}(st, \{(uid_i, \text{nonce}_i, \text{sig}'_i)\}_{i \in [n]}, e')$

uid_{lead} $\leftarrow \text{me}$

(e, p) $\leftarrow (st.e, st.p)$

$e_{\text{pub}} \leftarrow \perp$

$G \leftarrow \{uid_1, \dots, uid_n\} \cup \{\text{me}\}$

Added, Removed $\leftarrow \emptyset$

numLplLinks $\leftarrow \text{max-links}$

lpl $\leftarrow \text{*send-LPL}()$

hb $\leftarrow \text{*send-heartbeat}()$

M $\leftarrow (\text{me}, M_K, \text{lpl}, \text{hb})$

return M

Algorithm: Add(time, {(uid_i, sig_i)}_{i ∈ [n]})

req uid_{lead} = me

for $i \in [n]$ **do**

req uid_i $\notin G$

parse (sig'_i, nonce_i) $\leftarrow \text{sig}_i$

$G \leftarrow G \cup \{uid_1, \dots, uid_n\}$

Added $\leftarrow \text{Added} \cup \{uid_1, \dots, uid_n\}$

if $p \geq p_{\text{MAX}}$ **then**

try (st, M_K) $\leftarrow \text{cmKEM.Add}(st, \{(uid_i, \text{nonce}_i, \text{sig}'_i)\}_{i \in [n]}, \text{true})$

(e, p) $\leftarrow (st.e, st.p)$

else

try (st, M_K) $\leftarrow \text{cmKEM.Add}(st, \{(uid_i, \text{nonce}_i, \text{sig}'_i)\}_{i \in [n]}, \text{false})$

return (me, M_K, \perp , \perp)

Fig. 5: The proposed changes with respect to Zoom's scheme from Fig. 3.

The protocol. Concretely, our enhancement to Zoom's protocol is as follows: First, each party generates an unpredictable nonce nonce (from some nonce space \mathcal{N} , e.g., 192-bit strings) with frequency Δ_{nonce} . These nonces are seen as part of a party's credential and hence ParticipantTick outputs a new credential whenever the nonce is updated. (In practice one would of course only upload the nonce, not the entire credential, each time.)

Whenever a new leader is elected, they get each participants' latest nonce from the server. We encode this as part of the credentials sig_i for the Lead algorithm, which can now be thought as a sort of time-based credentials. The new leader then uses those nonces as associated data for the cmKEM primitive (which for Zoom's instantiation means it is used as associated data for the authenticated PKE). The same mechanism is used for adding new members to the group.

Each party as part of the Follow algorithm provides their current nonce as associated data to JoinSession, thus verifying that the new leader used the correct one. To prevent race conditions, parties moreover store their second latest nonce `nonce'` and try with that one if JoinSession initially fails. See Fig. 5 for a formal description of the changes with respect to Zoom’s current scheme from Fig. 3.

Security. Our proposal improves the liveness properties twofold. First, the liveness slack no longer degrades in the number of leader changes. Second, liveness now holds even if a *past leader* has been corrupted as long as the current leader is honest.

We now state the resulting theorem. A more formal version thereof and a proof can be found in Appendix E.1.

Theorem 3. *The modified LL-CGKA scheme from Fig. 5 is secure with the liveness slack of P being at most*

$$\min(t_{\text{now}} - t_{\text{joined}}, 2 \cdot \Delta_{\text{nonce}} + \Delta_{\text{live}}) + \Delta_{\text{live}},$$

where t_{now} denotes the current time, t_{joined} the time P joined the meeting. In contrast to Theorem 2, liveness holds if the current leader is honest (as opposed to all leaders encountered so far), analogous to all other properties. Additionally, as long as all n leaders encountered so far have been honest, the liveness slack of P is also at most $(n + 1) \cdot \Delta_{\text{live}}$.

Remark 2 (Zoom’s adaptation.). Zoom implemented a variant of this proposal in the Zoom client since version 5.13. The most significant difference is that, while our analysis treats the nonce interval Δ_{nonce} as a constant, in Zoom’s implementation for efficiency reasons Δ_{nonce} scales quadratically with the number of current meeting participants. It is easy to see that in our proposed protocol sending fresh nonces more frequently only improves security — hence Theorem 3 can be seen as worst-case analysis with Δ_{nonce} set to the maximal encountered during a meeting.

4.3 Leveraging Clock Synchronicity

In this section, we explore an alternative approach towards mitigating the degradation of liveness properties during leader changes. Concretely, we propose to leverage pre-existing clock synchronicity to achieve better liveness without having to introduce additional communication. For E2EE protocols, however, it is undesirable to simply assume synchronized clocks since this, for all practical purposes, implies assuming a trusted reference clock (such as a time server) and introduces additional friction for users on misconfigured devices (for example, with the wrong timezone settings.).

Unfortunately, even detecting whether clocks are synchronized is non-trivial. For instance, consider the interaction between a participant P and a leader L depicted in Fig. 6: in one situation, L ’s clock is in sync while in the other situation L ’s clock runs ahead — yet the scenarios look completely indistinguishable to both P and L . As such, we propose the following hybrid strategy:

- For correctness, i.e., functionality of the scheme, assume clocks to be properly synchronized. After all, Zoom is usually run on modern devices such as laptop computers or smartphones that generally do have well synchronized clocks. An honest Zoom server could moreover detect erroneous time settings and instruct the client to re-synchronize their clock (either by displaying a warning, or by doing it automatically with a somewhat trusted external server).
- For security, well-synchronized clocks should yield tight liveness assurances, while worst-case liveness should degrade to the current¹ protocol’s properties.

The protocol. We now discuss our proposed mechanism. In a nutshell, our proposed improvement works by each party P maintaining not just an upper bound $\delta_P[L]$ on the clock drift with their current leader L (which in this protocol we rename as $\delta_P^{\text{max}}[L]$ for clarity), but also a lower bound $\delta_P^{\text{min}}[L]$, such that $\delta_P^{\text{min}}[L] \leq \text{offset}_{L \rightarrow P} \leq \delta_P^{\text{max}}[L]$. As in the current protocol¹, these bounds are derived from simple causality observations, can gradually improve over the course of the execution, and in turn are used to adjust the timestamp indicated as part of the heartbeat messages and therefore decide when to drop out. See Fig. 7 for a formal description of the proposed modifications with respect to Zoom’s current scheme.



Fig. 6: The leader L 's clock running ahead (right) negatively affects liveness as the addition of P ' can be withheld longer from P .

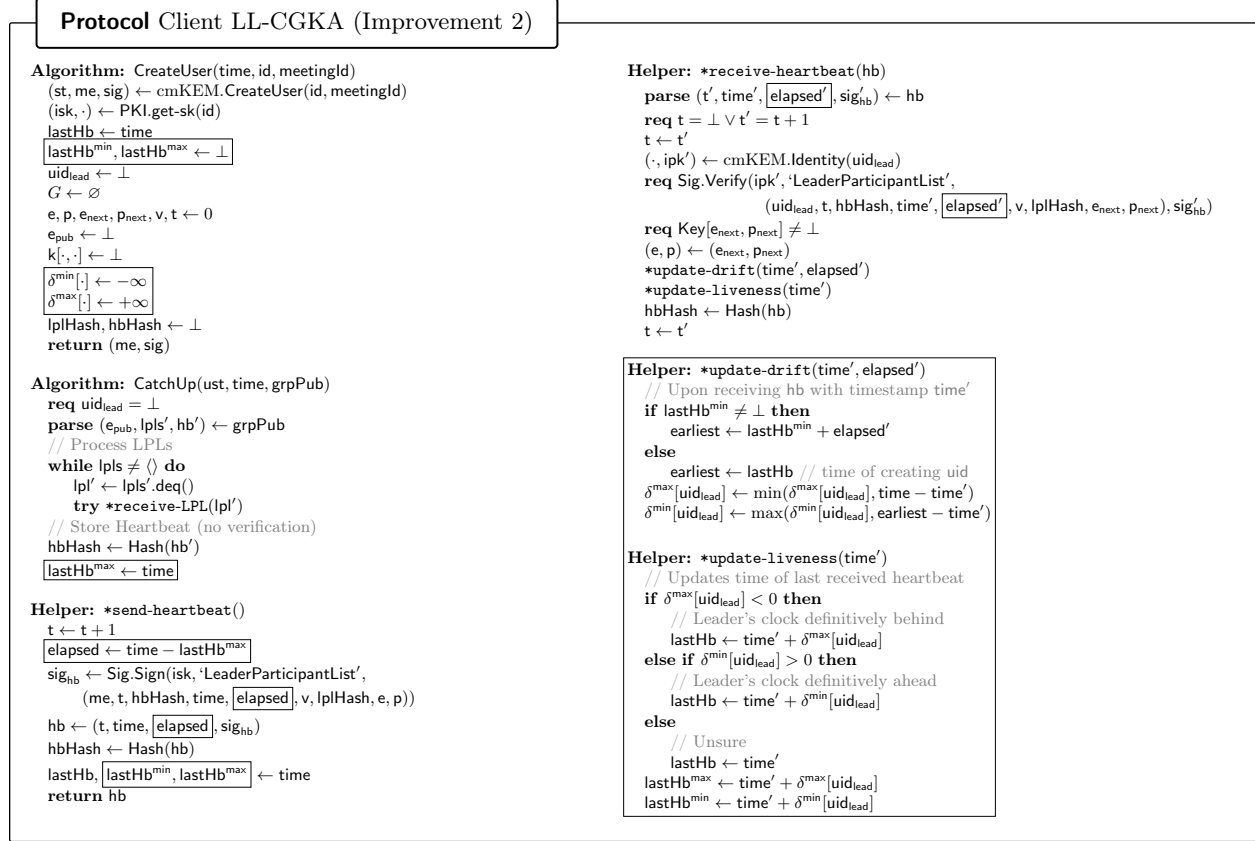


Fig. 7: The proposed changes with respect to Zoom's scheme from Fig. 3.

Deriving bounds. To this end, consider the case that P receives a heartbeat with timestamp t_L (according to L 's clock) at time t_{now} (according to P 's clock). As in Zoom's current protocol, P clearly knows that the heartbeat has not been sent after t_{now} , i.e. $t_L + \text{offset}_{L \rightarrow P} \leq t_{\text{now}}$. Furthermore, assume that (for whatever reason) P knows that this heartbeat has been sent definitively not before t_{earliest} . P can use this to deduce the following bounds:

$$t_{\text{earliest}} - t_L \leq \delta_P^{\min}[L] \quad \text{and} \quad \delta_P^{\max}[L] \leq t_{\text{now}} - t_L.$$

P will only update a bound if it improves the current one. (At the beginning, the protocol initializes them to $\delta_P^{\min}[L] = -\infty$ and $\delta_P^{\max}[L] = +\infty$.)

In our protocol, P will have a meaningful such lower bound t_{earliest} in the following two situations:

- **Upon joining the meeting:** When P joins the meeting, the first heartbeat they get will sign over an LPL containing their freshly generated ephemeral key. Hence, that heartbeat must have been sent after the time t_{joined} when P generated the key.
- **Upon receiving the first heartbeat from a new leader L' :** The protocol works by having P deduce a lower bound on when the last heartbeat from the old leader was sent, and the new leader L' indicating as part of the heartbeat a lower bound *elapsed* on the *elapsed duration* between the last heartbeat of the old leader L and their first one. Hence, upon receiving the first heartbeat from L' , P can use $\text{time}_L + \delta_P^{\min}[L] + \text{elapsed}$ as a lower bound on the sending time. Observe that the new leader L' can compute a lower bound *elapsed* based on the last heartbeat from L as follows: If L' has already been part of the meeting, it can leverage their own bound $\delta_{L'}^{\max}[L]$ to deduce the upper bound $\text{time}_L + \delta_{L'}^{\max}[L]$ on the prior heartbeat’s sending time. Otherwise, L' can use the time they got the last heartbeat from the server as part of **CatchUp** yielding at least some (very conservative) bound.

For subsequent heartbeats of the same leader, P only updates the upper bound (if tighter than the previous one).

Correcting the drift. We then modify the “conversion” of timestamp that P performs accordingly. That is, whenever P receives a heartbeat with timestamp time'_L , in Zoom’s protocol P knows that this has been sent no later than $\text{time}'_P := \text{time}'_L + \delta_P[L]$ and conservatively delays dropping out until $\text{time}'_P + \Delta_{\text{live}}$. Unfortunately, after a number of leader changes the tightness of the bound $\delta_P[L]$ degrades (i.e. the difference between $\delta_P[L]$ and the actual $\text{offset}_{L \rightarrow P}$ can become quite large). In other words, Zoom’s protocol favors correctness over liveness. Instead, this improved protocol adjusts the received timestamp if and only if the leader’s clock is surely behind or ahead, respectively:

$$\text{time}'_P := \begin{cases} \text{time}'_L + \delta_P^{\min}[L] & \text{if } \delta_P^{\min}[L] > 0, \\ \text{time}'_L + \delta_P^{\max}[L] & \text{if } \delta_P^{\max}[L] < 0, \\ \text{time}'_L & \text{otherwise.} \end{cases}$$

This results in each participant potentially dropping out earlier than they would in Zoom’s protocol (as the drift adjustment here is bounded by the one in Zoom’s protocol), allowing us to prove tighter liveness guarantees. At the same time, if clocks are synchronized the protocol will not make any drift adjustments and thus participants won’t prematurely drop out, ensuring a smooth meeting experience (i.e. provable correctness conditions, formalized in Appendix E.4).

Security. We now quantify the strengthened liveness properties of this scheme. A more formal version thereof and a proof is given in Appendix E.3.

Theorem 4. *The modified LL-CGKA scheme from Fig. 7 is secure with the following improved liveness slack*

$$\min(|\text{offset}_{L \rightarrow P}|, n \cdot \Delta_{\text{live}}, t_{\text{now}} - t_{\text{joined}}) + \Delta_{\text{live}},$$

where $\text{offset}_{L \rightarrow P}$ denotes the clock drift between P and their respective leader L , t_{now} denotes the current time, t_{joined} the time P joined the meeting, and n denotes the number of distinct leaders P encountered so far. Liveness holds if all those leaders have followed the protocol, while all other properties hold as long as the current leader is honest.

5 Meeting Stream Security

The notion of LL-CGKA formalizes the key agreement portion of Zoom’s E2EE meeting protocol. While our formal analysis stops at the level of the key agreement, we now comment on how these guarantees extend to the full protocol.

The symmetric meeting key that participants agree upon is leveraged in a straightforward way to provide security guarantees for the whole meeting, by composing it with AEAD. Concretely, given the meeting key, Zoom clients derive a specific per-stream subkey by using HKDF and mixing in a specific stream identifier which depends on the stream type as well as the participant identifier. This subkey is used by each participant to encrypt their streams using AES-GCM. Incrementing nonces provide protection against replay and out of order delivery.

Confidentiality and authenticity. Informally, confidentiality of the meeting key (as formalized in the LL-CGKA abstraction) implies confidentiality of the streams, as distinguishing encrypted meeting streams from encryptions of random noise would require breaking the AEAD scheme. Similarly, AEAD provides integrity protection against external attackers who do not have access to the meeting key, guaranteeing that any received ciphertexts was produced by someone with knowledge of the meeting (sub)key. As noted in the whitepaper [11], it is possible for attendees with privileged network access to tamper with each other’s streams.

Liveness. The liveness properties proven for the LL-CGKA directly guarantee that group operations in an E2EE meeting cannot be withheld, and extend analogously to the encrypted meeting streams, but with different parameters. Indeed, as of version 5.13 of the Zoom client, meeting participants stop decryption using old meeting keys shortly after a newer one is advertised from the key agreement, i.e., the LL-CGKA scheme (with a tolerance $\Delta_{\text{stream}} = 10$ seconds to account for network latency). In addition, meeting leaders rotate these keys at least once every $t = 5$ minutes even when there is no change in the participant list. Assuming the above, the protocol guarantees that each packet sent by an honest participant and successfully decrypted was sent within $t + \Delta_{\text{stream}} + \Delta$ of its decryption, where Δ is the liveness slack from the key agreement protocol. Alternatively, the protocol could include the heartbeat counter from the key agreement as associated data in the video encryption, yielding liveness $\Delta + \Delta_{\text{stream}} + \Delta_{\text{heartbeat}}$ without the need to frequently re-key.¹²

6 Conclusions

In this work, we provided the first formal security analysis of Zoom’s E2EE meetings protocol, which is one of the most popular group video communication tools in the world. Our work lead to a deployed improvement of the Zoom E2EE meetings protocol, which strengthens its security properties. Of independent interest, our work is also the first that defines and studies *liveness* in the context of end-to-end encryption, which we hope should find other applications beyond Zoom meetings.

7 Acknowledgements

We thank Michael Maxim and the rest of the Zoom team for review and implementation of the end-to-end encrypted meeting protocol and liveness improvements detailed in this work.

References

1. Alwen, J., Blanchet, B., Hauck, E., Kiltz, E., Lipp, B., Riepel, D.: Analysing the HPKE standard. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 87–116. Springer, Heidelberg (Oct 2021). https://doi.org/10.1007/978-3-030-77870-5_4
2. Alwen, J., Coretti, S., Dodis, Y.: The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 129–158. Springer, Heidelberg (May 2019). https://doi.org/10.1007/978-3-030-17653-2_5
3. Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Security analysis and improvements for the IETF MLS standard for group messaging. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part I. LNCS, vol. 12170, pp. 248–277. Springer, Heidelberg (Aug 2020). https://doi.org/10.1007/978-3-030-56784-2_9

¹² This is, however, non-trivial to achieve in a backwards compatible way.

4. Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Modular design of secure group messaging protocols and the security of MLS. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021. pp. 1463–1483. ACM Press (Nov 2021). <https://doi.org/10.1145/3460120.3484820>
5. Alwen, J., Coretti, S., Jost, D., Mularczyk, M.: Continuous group key agreement with active security. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part II. LNCS, vol. 12551, pp. 261–290. Springer, Heidelberg (Nov 2020). https://doi.org/10.1007/978-3-030-64378-2_10
6. An, J.H.: Authenticated encryption in the public-key setting: Security notions and analyses. Cryptology ePrint Archive, Report 2001/079 (2001), <https://eprint.iacr.org/2001/079>
7. Apple: Facetime & privacy. <https://www.apple.com/legal/privacy/data/en/face-time/>
8. Barnes, R., Beurdouche, B., Millican, J., Omara, E., Cohn-Gordon, K., Robert, R.: The messaging layer security (mls) protocol (draft-ietf-mls-protocol-latest). Tech. rep., IETF (Oct 2020), <https://messaginglayersecurity.rocks/mls-protocol/draft-ietf-mls-protocol.html>
9. Bellare, M., Goldwasser, S.: Verifiable partial key escrow. In: Graveman, R., Janson, P.A., Neuman, C., Gong, L. (eds.) ACM CCS 97. pp. 78–91. ACM Press (Apr 1997). <https://doi.org/10.1145/266420.266439>
10. Bienstock, A., Fairoze, J., Garg, S., Mukherjee, P., Raghuraman, S.: What is the exact security of the signal protocol? Preprint (2021), https://cs.nyu.edu/~afb383/publication/uc_signal/uc_signal.pdf
11. Blum, J., Booth, S., Chen, B., Gal, O., Krohn, M., Len, J., Lyons, K., Marcedone, A., Maxim, M., Mou, M.E., Namavari, A., O’Connor, J., Rien, S., Steele, M., Green, M., Kissner, L., Stamos, A.: Zoom cryptography whitepaper – v4.0. https://github.com/zoom/zoom-e2e-whitepaper/raw/master/archive/zoom_e2e_v4.pdf (2022)
12. Boneh, D., Naor, M.: Timed commitments. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 236–254. Springer, Heidelberg (Aug 2000). https://doi.org/10.1007/3-540-44598-6_15
13. Brendel, J., Cremers, C., Jackson, D., Zhao, M.: The provable security of ed25519: Theory and practice. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 1659–1676 (2021). <https://doi.org/10.1109/SP40001.2021.00042>
14. Bresson, E., Chevassut, O., Pointcheval, D.: Dynamic group Diffie-Hellman key exchange under standard assumptions. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 321–336. Springer, Heidelberg (Apr / May 2002). https://doi.org/10.1007/3-540-46035-7_21
15. Canetti, R., Garay, J., Itkis, G., Micciancio, D., Naor, M., Pinkas, B.: Multicast security: a taxonomy and some efficient constructions. In: IEEE INFOCOM ’99. Conference on Computer Communications. Proceedings. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. The Future is Now (Cat. No.99CH36320). vol. 2, pp. 708–716 (1999)
16. Cisco: Zero-trust security for webex – white paper. <https://www.cisco.com/c/en/us/solutions/collateral/collaboration/white-paper-c11-744553.html> (2021)
17. Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. Journal of Cryptology **33**(4), 1914–1983 (Oct 2020). <https://doi.org/10.1007/s00145-020-09360-1>
18. Cohn-Gordon, K., Cremers, C.J.F., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. In: 2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017. pp. 451–466. IEEE (2017). <https://doi.org/10.1109/EuroSP.2017.27>, <https://doi.org/10.1109/EuroSP.2017.27>
19. Coron, J.S., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård revisited: How to construct a hash function. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 430–448. Springer, Heidelberg (Aug 2005). https://doi.org/10.1007/11535218_26
20. Denis, F.: The sodium cryptography library. <https://download.libsodium.org/doc/> (Jun 2013)
21. Dolev, D., Strong, H.R.: Polynomial algorithms for multiple processor agreement. In: 14th ACM STOC. pp. 401–407. ACM Press (May 1982). <https://doi.org/10.1145/800070.802215>
22. Dwork, C., Naor, M.: Pricing via processing or combatting junk mail. In: Brickell, E.F. (ed.) CRYPTO’92. LNCS, vol. 740, pp. 139–147. Springer, Heidelberg (Aug 1993). https://doi.org/10.1007/3-540-48071-4_10
23. Dwork, C., Naor, M., Sahai, A.: Concurrent zero-knowledge. In: 30th ACM STOC. pp. 409–418. ACM Press (May 1998). <https://doi.org/10.1145/276698.276853>
24. Feldman, P., Micali, S.: Optimal algorithms for byzantine agreement. In: 20th ACM STOC. pp. 148–161. ACM Press (May 1988). <https://doi.org/10.1145/62212.62225>
25. Fischlin, M., Günther, F.: Multi-stage key exchange and the case of Google’s QUIC protocol. In: Ahn, G.J., Yung, M., Li, N. (eds.) ACM CCS 2014. pp. 1193–1204. ACM Press (Nov 2014). <https://doi.org/10.1145/2660267.2660308>
26. Garay, J.A., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol with chains of variable difficulty. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 291–323. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63688-7_10

27. Gruszczyk, J.: End-to-end encryption for one-to-one microsoft teams calls now generally available. Microsoft Teams Blog – December 14, 2021. <https://techcommunity.microsoft.com/t5/microsoft-teams-blog/end-to-end-encryption-for-one-to-one-microsoft-teams-calls-now/ba-p/3037697> (12 2021)
28. Harder, E.J., Wallner, D.M.: Key Management for Multicast: Issues and Architectures. RFC 2627 (Jun 1999). <https://doi.org/10.17487/RFC2627>, <https://www.rfc-editor.org/info/rfc2627>
29. Isobe, T., Ito, R.: Security analysis of end-to-end encryption for zoom meetings. In: Baek, J., Ruj, S. (eds.) Information Security and Privacy. pp. 234–253. Springer International Publishing, Cham (2021)
30. Katz, J.: Efficient and non-malleable proofs of plaintext knowledge and applications. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 211–228. Springer, Heidelberg (May 2003). https://doi.org/10.1007/3-540-39200-9_13
31. Kim, Y., Perrig, A., Tsudik, G.: Tree-based group key agreement. Cryptology ePrint Archive, Report 2002/009 (2002), <https://eprint.iacr.org/2002/009>
32. Krawczyk, H.: Cryptographic extraction and key derivation: The HKDF scheme. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 631–648. Springer, Heidelberg (Aug 2010). https://doi.org/10.1007/978-3-642-14623-7_34
33. Krohn, M.: Zoom rolling out end-to-end encryption offering. Zoom Blog – October 14, 2020. <https://blog.zoom.us/zoom-rolling-out-end-to-end-encryption-offering/> (10 2020)
34. Lowe, G.: A hierarchy of authentication specifications. In: Proceedings 10th Computer Security Foundations Workshop. pp. 31–43 (1997). <https://doi.org/10.1109/CSFW.1997.596782>
35. Marlinspike, M., Perrin, T.: The double ratchet algorithm (11 2016), <https://whispersystems.org/docs/specifications/doublerratchet/doublerratchet.pdf>
36. Panjwani, S.: Tackling adaptive corruptions in multicast encryption protocols. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 21–40. Springer, Heidelberg (Feb 2007). https://doi.org/10.1007/978-3-540-70936-7_2
37. Pass, R., Seeman, L., Shelat, A.: Analysis of the blockchain protocol in asynchronous networks. In: Coron, J.S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part II. LNCS, vol. 10211, pp. 643–673. Springer, Heidelberg (Apr / May 2017). https://doi.org/10.1007/978-3-319-56614-6_22
38. Perrig, A., Song, D., Canetti, R., Tygar, J.D., Briscoe, B.: Timed Efficient Stream Loss-Tolerant Authentication (TESLA): Multicast Source Authentication Transform Introduction. IETF RFC 4082 (Informational) (2005)
39. Pinto, A., Poettering, B., Schuldt, J.C.: Multi-recipient encryption, revisited. p. 229–238. ASIA CCS '14, Association for Computing Machinery, New York, NY, USA (2014). <https://doi.org/10.1145/2590296.2590329>, <https://doi.org/10.1145/2590296.2590329>
40. Poettering, B., Rösler, P., Schwenk, J., Stebila, D.: SoK: Game-based security models for group key exchange. In: Paterson, K.G. (ed.) CT-RSA 2021. LNCS, vol. 12704, pp. 148–176. Springer, Heidelberg (May 2021). https://doi.org/10.1007/978-3-030-75539-3_7
41. Rivest, R.L., Shamir, A., Wagner, D.A.: Time-lock puzzles and timed-release crypto (1996)
42. Seggelmann, R., Tuexen, M., Williams, M.: Transport layer security (tls) and datagram transport layer security (dtls) heartbeat extension. IETF RFC 6520 (Standards Track) (2012)
43. Smart, N.P.: Efficient key encapsulation to multiple parties. In: Blundo, C., Cimato, S. (eds.) SCN 04. LNCS, vol. 3352, pp. 208–219. Springer, Heidelberg (Sep 2005). https://doi.org/10.1007/978-3-540-30598-9_15
44. WhatsApp: Whatsapp encryption overview (2017), retrieved 05/2020 from <https://www.whatsapp.com/security/WhatsApp-Security-Whitepaper.pdf>
45. Wire Swiss GmbH: Wire security whitepaper. <https://wire-docs.wire.com/download/Wire+Security+Whitepaper.pdf> (2021)
46. Yang, Z.: On constructing practical multi-recipient key-encapsulation with short ciphertext and public key. Sec. and Commun. Netw. **8**(18), 4191–4202 (dec 2015). <https://doi.org/10.1002/sec.1334>, <https://doi.org/10.1002/sec.1334>
47. Yuan, E.S.: Zoom acquires keybase and announces goal of developing the most broadly used enterprise end-to-end encryption offering. Zoom Blog – May 7, 2020. <https://blog.zoom.us/zoom-acquires-keybase-and-announces-goal-of-developing-the-most-broadly-used-enterprise-end-to-end-encryption-offering/> (5 2020)

A Preliminaries

A.1 Notation

For $\mathbb{N} := \{1, 2, \dots\}$ and $x \in \mathbb{N}$, we write $[x] := \{1, 2, \dots, x\}$. We use the notation $\{x_i\}_{i \in [n]}$ to denote the set $\{x_1, x_2, \dots, x_n\}$ and in slight abuse of notation x_j to denote the corresponding element if an ordering is clear from the context. We write $x \leftarrow a$ to assign the value a to the variable x , and for a set \mathcal{B} , $x \leftarrow_{\$} \mathcal{B}$ for sampling an element in \mathcal{B} uniformly at random. For a cyclic group \mathbb{G} , we use $\langle g \rangle$ to denote the subgroup generated by g and, thus, $\langle g \rangle = \mathbb{G}$ to denote that g is a generator. We use multiplicative group notation. The security parameter is denoted by κ .

A.2 Pseudocode and Games

We use pseudocode to describe both protocols and security games. Most algorithms we consider are stateful, i.e., they take a state as part of the input and produce an updated state as part of the output. In some cases, for simplicity, we omit this state from the pseudocode description.

Typically, those algorithms are fallible and in the case the given inputs are invalid for the current state, the algorithm can return the special error value \perp instead. We may call such algorithms with the keyword **try** prepended, resulting in the error to propagate to the calling routine, by unwinding all state changes to the calling routine and then returning \perp as well. We use the keyword **parse** to fallibly parse a message as a tuple of values, resulting in \perp if the message was malformed. Furthermore, we use the keyword **req** followed by a boolean condition to specify preconditions, i.e., terminate the calling procedure with \perp if the condition evaluates to **false**. In contrast, in games, the keyword **assert** followed by a boolean condition is used to denote winning conditions for the adversary with the special variable `won` of the game being set to **true** whenever the condition evaluates to **false**. The keyword **pub** followed by a variable denotes that the adversary has read-only access to that variable during the game.

Data structures. We make use of associative arrays. We denote by $A[i]$ the value at position i and by $A[i] \leftarrow x$ the respective assignment. Moreover, we use the special symbol \perp as a shorthand for uninitialized values — i.e., the condition $A[i] = \perp$ evaluates to **true** iff position i is not set and $A[i] \leftarrow \perp$ clears said position. We use $A[\cdot] \leftarrow \perp$ to initialize an empty array and $B[\cdot] \leftarrow x$ to initialize one with the value x at every position. Additionally, we make use of FIFO queues. We write $\langle \rangle$ to denote the empty queue, $Q.\text{enq}(x)$ for enqueueing x and $y \leftarrow Q.\text{deq}()$ for dequeueing an item and storing it in y . We moreover use $y \leftarrow Q.\text{peek}()$ to retrieve the first item without actually removing it from Q and $Q.\text{reverse}()$ to reverse the order of the queue.

Games. We write our games using the following conventions. Most games have two special algorithms, `Initialize` and `Finalize`, where the adversary \mathcal{A} initiates the interaction calling the former, is then allowed to make an arbitrary number of queries to all other oracles before ending the interaction with a single invocation to `Finalize`. The output bit of the oracle is then seen as the output of the interaction. Pure distinguishing games are written as two worlds representing the $b = 0$ and $b = 1$ experiments.

A.3 Cryptographic Primitives

Symmetric Encryption. A symmetric encryption scheme is a tuple of algorithms $\text{SE} := (\text{SE.Enc}, \text{SE.Dec})$. The encryption algorithm $\text{SE.Enc}: \text{SE.K} \times \text{SE.M} \rightarrow \{0, 1\}^*$ takes a key $k \in \text{SE.K}$ and a message $m \in \text{SE.M}$ to produce a ciphertext c . The deterministic decryption algorithm $\text{SE.Dec}: \text{SE.K} \times \{0, 1\}^* \rightarrow \text{SE.M} \cup \{\perp\}$, given the key and the ciphertext, outputs either a message $m \in \text{SE.M}$ or \perp .

For correctness, we require that $\Pr[\text{SE.Dec}(k, \text{SE.Enc}(k, m)) = m] = 1$, where the randomness is taken over the choice of $k \in \text{SE.K}$ and the encryption algorithm. For security, we either require the standard IND-CPA or IND-CCA2 notions, depending on the context.

Nonce-Based AEAD. A nonce-based authenticated encryption scheme with associated data is a tuple of deterministic algorithms $(\text{AEAD.Enc}, \text{AEAD.Dec})$. The algorithm $\text{AEAD.Enc}: \text{AEAD.K} \times \text{AEAD.N} \times \text{AEAD.M} \times \text{AEAD.AD} \rightarrow \{0, 1\}^*$ takes a key $k \in \text{AEAD.K}$ and a nonce $\text{nonce} \in \text{AEAD.N}$, a message $m \in \text{AEAD.M}$, and associated data $\text{ad} \in \text{AEAD.AD}$, to produce a ciphertext c . The decryption algorithm $\text{AEAD.Dec}: \text{AEAD.K} \times \text{AEAD.N} \times \{0, 1\}^* \times \text{AEAD.AD} \rightarrow \text{AEAD.M} \cup \{\perp\}$, given the key, the nonce, the ciphertext, and the associated data, outputs either a message $m \in \text{AEAD.M}$ or \perp .

For correctness, it is required that

$$\text{AEAD.Dec}(k, \text{nonce}, \text{AEAD.Enc}(k, \text{nonce}, m, \text{ad}), \text{ad}) = m$$

for all $k \in \text{AEAD.K}$, $\text{nonce} \in \text{AEAD.N}$, $m \in \text{AEAD.M}$, and $\text{ad} \in \text{AEAD.AD}$. Moreover, for simplicity we assume a fixed length message space, i.e., that $|m_i| = |m_j|$ for all $m_i, m_j \in \text{AEAD.M}$.

Security is then formalized using the distinguishing game between a real-world and ideal-world system in Fig. 8. The game simultaneously captures confidentiality (in the ideal-world a random message gets encrypted) and authenticity (the decryption of mauled/injected messages in the real-world must return \perp).

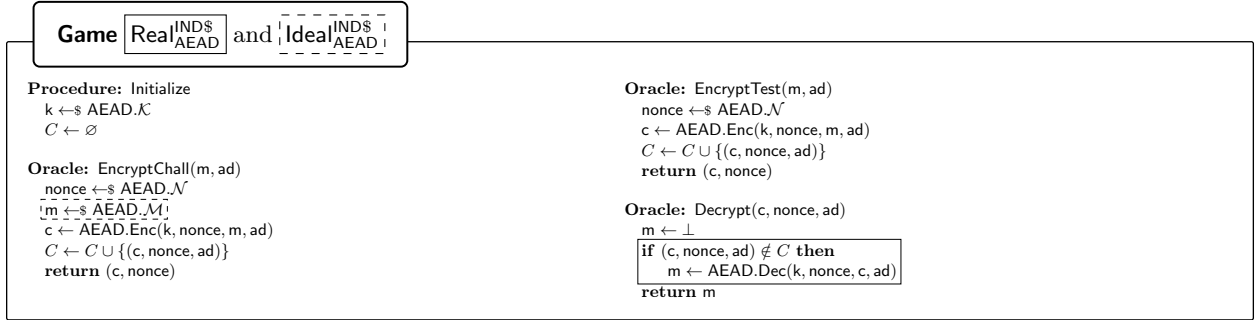


Fig. 8: The nonce-based AEAD security experiments.

Public-Key Encryption. We use a public-key encryption scheme $\text{PKE} := (\text{PKE.kg}, \text{PKE.enc}, \text{PKE.dec})$, where $(\text{sk}, \text{pk}) \leftarrow \text{PKE.kg}(1^\kappa)$ denotes the key generation, $c \leftarrow \text{PKE.enc}(\text{pk}, m)$ the encryption, and $m \leftarrow \text{PKE.dec}(\text{sk}, c)$ the decryption, respectively. For security, we require the standard IND-CCA2 notion.

Digital Signatures. Many of our constructions use a digital signature scheme $\text{Sig} := (\text{Sig.KeyGen}, \text{Sig.Sign}, \text{Sig.Verify})$. Note that for convenience we require both signing and verification to split the input into the actual message m and an additional context string context (used for domain separation when the signing key can be used for multiple purposes). Hence, for $(\text{isk}, \text{ipk}) \leftarrow \text{Sig.KeyGen}(1^\kappa)$ one signs a message under a context using $\text{sig} \leftarrow \text{Sig.Sign}(\text{isk}, \text{context}, m)$ and verifies the respective signature using $b \leftarrow \text{Sig.Verify}(\text{ipk}, \text{context}, m, \text{sig})$. For security, we require EUF-CMA security such that $b = 1$ only if both the message m and the context match.

HKDF. We use a key derivation function $\text{HKDF}: \mathbb{G} \times \{0, 1\}^* \rightarrow \text{AEAD.K}$, which we model as a random oracle. The function takes as input a high entropy string (in our case, an element g in the Diffie Hellman group from which keys are sampled) and a domain separation string context , and outputs a pseudorandom string of the appropriate length to be used as an AEAD key. To prove our schemes secure, we assume that this function behaves like a random oracle. In practice, Zoom’s protocol uses the HMAC based Key Derivation Function defined in RFC5689 [32] (using 0^λ as the salt for the HKDF.Expand function).

Pseudo Random Generator (PRG). For the cmKEM construction we require a stateful PRG scheme $\text{PRG} := (\text{PRG.Init}, \text{PRG.Eval})$. The algorithm $\text{seed} \leftarrow \text{PRG.Init}(1^\kappa)$ initializes a state seed , while $(\text{seed}', k) \leftarrow \text{PRG.Eval}(\text{seed})$ evaluates the PRG outputting a string k as well as an updated state seed' .

For security, we require that the sequence of outputs k_1, k_2, \dots is indistinguishable from a sequence of independent and uniformly random values (of the appropriate length).

A.4 Gap Diffie Hellman

The Zoom scheme makes use a cyclic group $\mathbb{G} = \langle g \rangle$ with a fixed generator g (technically a family of groups indexed by the security parameter) for which the Gap Diffie-Hellman (Gap-DH) is assumed to be hard. Gap-DH states that the Computational Diffie-Hellman (CDH) problem remains hard even when given access to a decisional oracle taking $(X, Y, Z) \in \mathbb{G}^3$ and returning 1 iff $\text{DH}(X, Y) = Z$.

We remark that the Gap-DH assumption (rather than, e.g., CDH) appears to be rather intrinsic to this kind of simple Diffie-Hellman based protocol, as exemplified by recent analyses of Signal [10] or the HPKE standard [1].

B Zoom’s PKI

In analysis, we assume a simple (long-term) public-key infrastructure to make formal statements that assure security properties assuming user authenticity. The PKI provides each long-term identity id with their own private signing key isk , while allowing all other users to verify that the respective public verification key ipk belongs to id . Such a PKI is, however, not how — at the time of writing — Zoom actually verifies public keys, and we refer to the whitepaper [11] for further details on Zoom’s ongoing efforts for improving user authentication. In the following, we briefly summarize those efforts at the time of writing.

Security codes. Currently, each user does have long-term keys but there is no actual PKI. Instead, security relies on the use of the so-called *meeting leader security code*, which is a digest of the long-term public key. The meeting leader is supposed to read out their security code at the beginning of the meeting (or whenever a new leader takes over) and all participants compare the code to what their client displays. Assuming an attacker cannot convincingly forge audio and video data, and that all participants personally know the meeting leader, this ensures that participants only stay in the meeting if they have the right key. Moreover, the leader is assured of the participants’ identities by having them explicitly acknowledge the match of the security code. In a nutshell, this mechanism piggybacks on the meeting’s E2E encryption to leverage the unidirectional security, established by the security code, to bidirectional one. A bit more concretely, in the meeting the leader securely distributes a shared symmetric key to all participants that is then used to encrypt the video stream. Now the leader security code ensures that a MITM cannot tamper with this key distribution process and, hence, the key material is known exactly to the set of recipients addressed by the leader. (At this point, the recipients may or may not be the intended ones.) If now each of them speaks up, and the leader knows them personally, the leader can verify that the set of participants matches their expectation.

Performed correctly, and under the right set of assumptions, an attacker controlling Zoom’s infrastructure can, thus, not break long-term key authenticity.

Identity Provider Attestations. Zoom plans to introduce more sophisticated user identities which will contain additional data other than a changeable display name, such as email addresses and account identifiers. Organizations leveraging external Identity Providers (IDP) to manage authentication and access control for their members, will be able to delegate their IDP to attest those members’ identities using an extension of OpenID Connect. At this point, our simplistic PKI will become more accurate, with the assumption of the PKI essentially corresponding to assuming the external IDPs to be honest.

Key Transparency. Further, Zoom envisions deploying a PKI based on a transparency tree. This complements the guarantees from external IDPs, especially for end-users not belonging to an organization providing an IDP. While this key transparency infrastructure will be hosted by Zoom, its append-only and public verifiability properties will make misbehavior detectable, resulting in stronger security.

C Details on cmKEM

C.1 The PKI

Let us briefly formalize the PKI assumed by both the protocol and the respective security game. See Appendix B for a discussion of how this PKI model applies to Zoom. The PKI provides the following interface:

- A user id can fetch their key pair. That is, when calling $\text{PKI.get-sk}(id)$, the PKI looks up whether a triple (id, isk, ipk) is recorded. Otherwise, it samples a new key pair using $(isk, ipk) \leftarrow \text{Sig.KeyGen}(1^\kappa)$ and stores the respective triple. Finally, it returns (isk, ipk) .
- Any user id can verify the public key of another user id' . That is, when id calls $\text{PKI.verify-pk}(id', ipk')$, the PKI looks up whether a triple (id', isk', ipk') is recorded for some isk' and returns 1 iff so.

C.2 Correctness

Intuitively, correctness formalizes that all participants get the same sequence of keys. On one hand, this means ensuring that the invocations to the cmKEM algorithms succeed, i.e., do not abort when provided proper inputs, such as when the leader only adds existing parties who are not yet members of the group or when participants process honestly generated messages. On the other hand, it means ensuring that the keys output by all participants, as well as the leader, are indeed equal. (Note that the second property is formally implied by the consistency properties of the security game. For clarity, we nevertheless require it as part of correctness as well.)

Due to the rather complex interaction, correctness is formalized as a game depicted in Fig. 9. The game largely follows the structure of the security game (see Section 2.3 and Appendix C.3), with the main exception that participants immediately process their respective messages — as part of the $\text{*verifyParticipants}$ helper function — rather than the adversary being able to schedule the delivery.

C.3 Security

In this section, we present the formal security definition of a cmKEM scheme. The security game is depicted in Fig. 10. We refer to Section 2.3 for a high-level discussion of the security game and in the following outline some of the more technical aspects.

Basic game state. The game keeps track of each user’s current protocol state, leader, epoch, and period, using the arrays St , Leader , Epoch , and Period , respectively. Moreover, for each session, the game also keeps track of the keys and rosters. It is thereby assumed that each state is uniquely identified by the triple $(\text{uid}_{\text{lead}}, e, p)$, i.e., the session’s leader as well as the epoch and period within a session. As a result, the values are stored in $\text{Key}[\text{uid}_{\text{lead}}, e, p]$ and $\text{Group}[\text{uid}_{\text{lead}}, e, p]$, respectively.

The St and Leader are updated directly on the fly. The epoch and period are updated as part of the *verifyProgress helper method, which first checks that the updated values, as output by the protocol, match the expected ones. Similarly, Key is updated as part of the $\text{*verifyConsistency}$ method that first checks key consistency. Group is updated via the *setGroup helper method. Finally, the game keeps track of corruptions and challenges as explained below.

Corruptions and member authentication. The game keeps track of corruptions using (1) the CorrIds set and (2) the Corrupted array. The former keeps track of all the corrupted long-term identities, while for the latter $\text{Corrupted}[\text{uid}]$ stores whether uid has been corrupted. It is modeled as an array (rather than set) for the following technicality: each uid starts out to be corrupted at the beginning of the game, then becomes “uncorrupted” once honestly generated, and may later become corrupted again. To this end, whenever the adversary chooses to corrupt id , then this also marks all associated ephemeral users that are *still active* as corrupted. (The game has a special DeleteUser oracle that allows a party to signify that it wants to terminate a ephemeral identity without having to be explicitly removed by another party.)

Game Corr _{Ψ, \mathcal{A}} cmKEM

Main

Procedure: Initialize
 won \leftarrow false
 pub \leftarrow Ψ .InitSplitState()
 pub St $[\cdot]$, Sigs $[\cdot]$, Epoch $[\cdot]$, Period $[\cdot]$, Leader $[\cdot]$,
 Group $[\cdot, \cdot, \cdot]$, Key $[\cdot, \cdot, \cdot]$ \leftarrow \perp

Procedure: Finalize
 return won

Users and session management

Oracle: CreateUser(id, meetingId)
 (st, uid, sig) \leftarrow Ψ .CreateUser(id, meetingId)
 (id', ipk) \leftarrow Ψ .Identity(uid)
 assert id = id' \wedge Ψ .Meeting(uid) = meetingId
 assert PKI.verify-pk(id, ipk)
 assert st \neq \perp \wedge St[uid] = \perp
 St[uid] \leftarrow st
 Sigs[uid] \leftarrow sig
 return uid

Oracle: StartSession(uid_{lead}, {(uid_i, ad_i, sig_i)_i \in [n]}, e')

req St[uid_{lead}] \neq \perp \wedge *isLeader(uid_{lead})
 req e' = \perp \vee Epoch[uid_{lead}] = \perp \vee e' > Epoch[uid_{lead}]
 req $\forall i \in [n] : \Psi$.Meeting(uid_i) = Ψ .Meeting(uid_{lead})
 \wedge uid_i \neq uid_{lead}
 req $\forall i, j \in [n], i < j : uid_i \neq uid_j$
 (st', M) \leftarrow Ψ .StartSession(St[uid_{lead}], {(uid_i, ad_i, sig_i)_i \in [n]}, e')

if (st', M) \neq \perp then
 St[uid_{lead}] \leftarrow st'
 Leader[uid_{lead}] \leftarrow uid_{lead}
 G' \leftarrow {(uid_i, ad_i)_i \in [n]} \cup {(uid_{lead}, \perp)}
 if e' \neq \perp then Epoch[uid_{lead}] \leftarrow e' - 1
 else Epoch[uid_{lead}] \leftarrow -1
 assert *verifyCorrectness(uid_{lead}, G', M, true)
 return M

else
 assert $\exists i \in [n] : St[uid_i] = \perp \vee Sigs[uid_i] \neq sig_i$
 return \perp

Group and key management

Oracle: Add(uid_{lead}, {(uid_i, ad_i, sig_i)_i \in [n]}, newEpoch)

req St[uid_{lead}] \neq \perp \wedge *isLeader(uid_{lead})
 (e, p) \leftarrow (Epoch[uid_{lead}], Period[uid_{lead}])
 req $\forall i \in [n] : (uid_i, *) \notin$ Group[uid_{lead}, e, p] \wedge uid_i \neq uid_{lead}
 \wedge Ψ .Meeting(uid_i) = Ψ .Meeting(uid_{lead})

req $\forall i, j \in [n], i < j : uid_i \neq uid_j$
 (st', M) \leftarrow Ψ .Add(St[uid_{lead}], {(uid_i, ad_i, sig_i)_i \in [n]}, newEpoch)

if (st', M) \neq \perp then
 St[uid_{lead}] \leftarrow st'
 G' \leftarrow Group[uid_{lead}, e, p] \cup {(uid_i, ad_i)_i \in [n]}
 assert *verifyCorrectness(uid_{lead}, G', M, newEpoch)
 return M

else
 assert $\exists i \in [n] : St[uid_i] = \perp \vee Sigs[uid_i] \neq sig_i$
 return \perp

Oracle: Remove(uid_{lead}, {uid_i}_i \in [n])

req St[uid_{lead}] \neq \perp \wedge *isLeader(uid_{lead})
 (e, p) \leftarrow (Epoch[uid_{lead}], Period[uid_{lead}])
 req uid_{lead} \notin {uid_i}_i \in [n] \wedge {(uid_i, *)_i \in [n]} \subseteq Group[uid_{lead}, e, p]
 (St[uid_{lead}], M) \leftarrow Ψ .Remove(St[uid_{lead}], {uid_i}_i \in [n])
 assert (St[uid_{lead}], M) \neq \perp
 G' \leftarrow Group[uid_{lead}, e, p] \setminus {(uid_i, *)_i \in [n]}
 assert *verifyCorrectness(uid_{lead}, G', M, true)
 return M

Helper Methods

Helper: *isLeader(uid)
 return Leader[uid] = uid

Helper: *verifyCorrectness(uid_{lead}, G', M, newEpoch)
 (e, p) \leftarrow (Epoch[uid_{lead}], Period[uid_{lead}])
 G \leftarrow Group[uid_{lead}, e, p]
 (e', p', k') \leftarrow (St[uid_{lead}], e, St[uid_{lead}].p, St[uid_{lead}].k')
 if \neg *verifyLeader(uid_{lead}, e, p, e', p', k', G', newEpoch) then
 return false
 if \neg *verifyParticipants(uid_{lead}, e', p', k', G', M) then
 return false
 return true

Helper: *verifyLeader(uid_{lead}, e, p, e', p', k', G', newEpoch)
 (Epoch[uid_{lead}], Period[uid_{lead}]) \leftarrow (e', p')
 (Group[uid_{lead}, e', p'], Key[uid_{lead}, e', p']) \leftarrow (G', k')
 if newEpoch then
 return e' = e + 1 \wedge p' = 0
 else
 return e' = e \wedge p' = p + 1

Helper: *verifyParticipants(uid_{lead}, e', p', k', G', M)
 (pub, ms) \leftarrow Split(pub, M)
 if pub.e \neq e' then return false
 for all uid : ms[uid] \neq \perp do
 if (uid, *) \notin G' then return false
 for all (uid, *) \in G' do
 if ms[uid] = \perp \wedge uid \neq uid_{lead} then return false
 for all uid : uid \neq uid_{lead} \wedge St[uid] \neq \perp \wedge (uid, *) \in G' do
 processed \leftarrow false
 if (uid, *) \in G' then
 St[uid] \leftarrow Ψ .Process(St[uid], ms[uid])
 processed \leftarrow true
 else if e' > Epoch[uid] then
 let ad s.t. (uid, ad) \in G'
 St[uid] \leftarrow Ψ .JoinSession(St[uid], uid_{lead}, Sigs[uid_{lead}], ms[uid], ad)
 Leader[uid] \leftarrow uid_{lead}
 processed \leftarrow true
 if processed then
 if St[uid] = \perp \vee e' \neq St[uid].e \vee p' \neq St[uid].p \vee k' \neq St[uid].k then
 return false
 (Epoch[uid], Period[uid]) \leftarrow (e', p')

Fig. 9: The correctness game for a cmKEM scheme Ψ .

Game $\text{Sec}_{\Psi, \mathcal{A}}^{\text{cmKEM}}$

Main

Procedure: Initialize

```

 $b \leftarrow \{0, 1\}$ 
 $\text{pub} \leftarrow \Psi.\text{InitSplitState}()$ 
 $\text{CorrIds}, \text{Challs} \leftarrow \emptyset$ 
 $\text{St}[\cdot], \text{Epoch}[\cdot], \text{Period}[\cdot], \text{Leader}[\cdot],$ 
    $\text{Group}[\cdot, \cdot, \cdot], \text{Key}[\cdot, \cdot, \cdot], \text{ChallKeys}[\cdot, \cdot, \cdot] \leftarrow \perp$ 
 $\text{Corrupted}[\cdot] \leftarrow \text{true}$ 
return pub

```

Procedure: Finalize(b')

```

if  $\neg * \text{safe}()$  then return false
else if won then return true
else return  $b' = b$ 

```

Session management

Oracle: CreateUser(id, meetingId)

```

 $(\text{st}, \text{uid}, \text{sig}) \leftarrow \Psi.\text{CreateUser}(\text{id}, \text{meetingId})$ 
 $(\text{id}', \text{ipk}) \leftarrow \Psi.\text{Identity}(\text{uid})$ 
assert  $\text{id} = \text{id}' \wedge \Psi.\text{Meeting}(\text{uid}) = \text{meetingId}$ 
assert  $\text{PKI.verify-pk}(\text{id}, \text{ipk})$ 
assert  $\text{St}[\text{uid}] = \perp$ 
 $\text{St}[\text{uid}] \leftarrow \text{st}$ 
 $\text{Corrupted}[\text{uid}] \leftarrow \text{false}$ 
return  $(\text{uid}, \text{sig})$ 

```

Oracle: StartSession($\text{uid}_{\text{lead}}, \{(\text{uid}_i, \text{ad}_i, \text{sig}_i)\}_{i \in [n]}, e'$)

```

req  $\text{St}[\text{uid}_{\text{lead}}] \neq \perp$ 
try  $(\text{St}[\text{uid}_{\text{lead}}], M) \leftarrow \Psi.\text{StartSession}(\text{St}[\text{uid}_{\text{lead}}], \{(\text{uid}_i, \text{ad}_i, \text{sig}_i)\}_{i \in [n]}, e')$ 
if  $e' \neq \perp$  then
  assert  $e' = \text{St}[\text{uid}_{\text{lead}}].e$ 
   $\text{Leader}[\text{uid}_{\text{lead}}] \leftarrow \text{uid}_{\text{lead}}$ 
  for  $i \in [n]$  do
    assert  $* \text{verifyCredentials}(\text{uid}_i, \text{sig}_i)$ 
  assert  $* \text{verifyProgress}(\text{uid}_{\text{lead}}, \text{'startedSession'})$ 
   $* \text{setGroup}(\text{uid}_{\text{lead}}, \{(\text{uid}_i, \text{ad}_i)\}_{i \in [n]} \cup \{(\text{uid}_{\text{lead}}, \perp)\})$ 
  assert  $* \text{verifyConsistency}(\text{uid}_{\text{lead}}, *)$ 
  return  $(M, * \text{pubState}(\text{uid}_{\text{lead}}))$ 

```

Oracle: JoinSession($\text{uid}, \text{uid}_{\text{lead}}, \text{sig}_{\text{lead}}, m, \text{ad}$)

```

req  $\text{St}[\text{uid}] \neq \perp \wedge \text{uid} \neq \text{uid}_{\text{lead}}$ 
try  $\text{St}[\text{uid}] \leftarrow \Psi.\text{JoinSession}(\text{St}[\text{uid}], \text{uid}_{\text{lead}}, \text{sig}_{\text{lead}}, m, \text{ad})$ 
 $\text{Leader}[\text{uid}] \leftarrow \text{uid}_{\text{lead}}$ 
assert  $* \text{verifyCredentials}(\text{uid}_{\text{lead}}, \text{sig}_{\text{lead}})$ 
assert  $* \text{verifyProgress}(\text{uid}, \text{'joinedSession'})$ 
assert  $* \text{verifyConsistency}(\text{uid}, \text{ad})$ 
return  $* \text{pubState}(\text{uid})$ 

```

Oracle: DeleteUser(uid)

```

 $\text{St}[\text{uid}] \leftarrow \perp$ 

```

Message processing (participants)

Oracle: Process(uid, m)

```

req  $\text{St}[\text{uid}] \neq \perp \wedge \neg * \text{isLeader}(\text{uid})$ 
try  $\text{St}[\text{uid}] \leftarrow \Psi.\text{Process}(\text{St}[\text{uid}], m)$ 
assert  $* \text{verifyProgress}(\text{uid}, \text{'eitherChanged'})$ 
assert  $* \text{verifyConsistency}(\text{uid}, *)$ 
return  $* \text{pubState}(\text{uid})$ 

```

Group and key management (leader)

Oracle: Add($\text{uid}_{\text{lead}}, \{(\text{uid}_i, \text{ad}_i, \text{sig}_i)\}_{i \in [n]}, \text{newEpoch}$)

```

req  $\text{St}[\text{uid}_{\text{lead}}] \neq \perp \wedge * \text{isLeader}(\text{uid}_{\text{lead}})$ 
 $(e, p) \leftarrow (\text{Epoch}[\text{uid}_{\text{lead}}], \text{Period}[\text{uid}_{\text{lead}}])$ 
try  $(\text{St}[\text{uid}_{\text{lead}}], M) \leftarrow \Psi.\text{Add}(\text{St}[\text{uid}_{\text{lead}}],$ 
    $\{(\text{uid}_i, \text{ad}_i, \text{sig}_i)\}_{i \in [n]}, \text{newEpoch})$ 
if newEpoch then
  assert  $* \text{verifyProgress}(\text{uid}_{\text{lead}}, \text{'epochChanged'})$ 
else
  assert  $* \text{verifyProgress}(\text{uid}_{\text{lead}}, \text{'periodChanged'})$ 
for all  $i \in [n]$  do
  assert  $(\text{uid}_i, *) \notin \text{Group}[\text{uid}_{\text{lead}}, e, p]$ 
    $\wedge \Psi.\text{Meeting}(\text{uid}_i) = \Psi.\text{Meeting}(\text{uid}_{\text{lead}})$ 
    $\wedge * \text{verifyCredentials}(\text{uid}_i, \text{sig}_i)$ 
 $* \text{setGroup}(\text{uid}_{\text{lead}}, \text{Group}[\text{uid}_{\text{lead}}, e, p] \cup \{(\text{uid}_i, \text{ad}_i)\}_{i \in [n]})$ 
assert  $* \text{verifyConsistency}(\text{uid}_{\text{lead}}, *)$ 
return  $(M, * \text{pubState}(\text{uid}_{\text{lead}}))$ 

```

Oracle: Remove($\text{uid}_{\text{lead}}, \{\text{uid}_i\}_{i \in [n]}$)

```

req  $\text{St}[\text{uid}_{\text{lead}}] \neq \perp \wedge * \text{isLeader}(\text{uid}_{\text{lead}})$ 
 $(e, p) \leftarrow (\text{Epoch}[\text{uid}_{\text{lead}}], \text{Period}[\text{uid}_{\text{lead}}])$ 
try  $(\text{St}[\text{uid}_{\text{lead}}], M) \leftarrow \Psi.\text{Remove}(\text{St}[\text{uid}_{\text{lead}}], \{\text{uid}_i\}_{i \in [n]})$ 
assert  $* \text{verifyProgress}(\text{uid}_{\text{lead}}, \text{'epochChanged'})$ 
assert  $\text{uid}_{\text{lead}} \notin \{\text{uid}_i\}_{i \in [n]} \wedge \{(\text{uid}_i, *)\}_{i \in [n]} \subseteq \text{Group}[\text{uid}_{\text{lead}}, e, p]$ 
 $* \text{setGroup}(\text{uid}_{\text{lead}}, \text{Group}[\text{uid}_{\text{lead}}, e, p] \setminus \{(\text{uid}_i, *)\}_{i \in [n]})$ 
assert  $* \text{verifyConsistency}(\text{uid}_{\text{lead}}, *)$ 
return  $(M, * \text{pubState}(\text{uid}_{\text{lead}}))$ 

```

Challenges & corruptions

Oracle: Test($\text{uid}_{\text{lead}}, e, p$)

```

req  $\text{Key}[\text{uid}_{\text{lead}}, e, p] \neq \perp$ 
if  $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p] = \perp$  then
   $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p] \leftarrow \text{Key}[\text{uid}_{\text{lead}}, e, p]$ 
return  $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p]$ 

```

Oracle: Challenge($\text{uid}_{\text{lead}}, e, p$)

```

req  $\text{Key}[\text{uid}_{\text{lead}}, e, p] \neq \perp$ 
 $\text{Challs} \leftarrow \text{Challs} \cup \{(\text{uid}_{\text{lead}}, e, p)\}$ 
if  $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p] = \perp$  then
  if  $b = 1$  then
     $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p] \leftarrow \text{Key}[\text{uid}_{\text{lead}}, e, p]$ 
  else
     $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p] \leftarrow \mathcal{K}$ 
return  $\text{ChallKeys}[\text{uid}_{\text{lead}}, e, p]$ 

```

Oracle: Corrupt(id)

```

 $\text{CorrIds} \leftarrow \text{CorrIds} \cup \{\text{id}\}$ 
 $L \leftarrow \emptyset$ 
for all  $\text{uid} : \text{St}[\text{uid}] \neq \perp$  do
  if  $\text{Identity}(\text{uid}) = (\text{id}, \cdot)$  then
     $\text{Corrupted}[\text{uid}] \leftarrow \text{true}$ 
     $L \leftarrow L \cup \{\text{St}[\text{uid}]\}$ 
 $(\text{isk}, \cdot) \leftarrow \text{PKI.get-sk}(\text{id})$ 
return  $(L, \text{isk})$ 

```

PKI Interaction

Oracle: Verify-Pk(id, ipk)

```

return  $\text{PKI.verify-pk}(\text{id}, \text{ipk})$ 

```

Oracle: Sign(id, context, m)

```

req  $\text{context} \neq \text{'EncryptionKeyAnnouncement'}$ 
 $(\text{isk}, \cdot) \leftarrow \text{PKI.get-sk}(\text{id})$ 
return  $\text{Sig.Sign}(\text{isk}, \text{context}, m)$ 

```

Fig. 10: The game formalizing the security of a cmKEM scheme Ψ .

Game Sec _{Ψ, \mathcal{A}} cmKEM Helpers

```

Helper: *pubState(uid)
  return (Epoch[uid], Period[uid])

Helper: *isLeader(uid)
  return Leader[uid] = uid

Helper: *members(uidlead, e, p)
  if Group[uidlead, e, p]  $\neq \perp$  then
    return Group[uidlead, e, p]
  else if p > 0 then
    return *members(uidlead, e, p - 1)
  else return  $\emptyset$ 

Helper: *verifyCredentials(uid)
  (id,  $\cdot$ )  $\leftarrow \Psi$ .Identity(uid)
  return id  $\in$  Corrlids  $\vee$   $\neg$ Corrupted[uid]

Helper: *verifyProgress(uid, expected)
  (e, p)  $\leftarrow$  (Epoch[uid], Period[uid])
  (e', p')  $\leftarrow$  (St[uid].e, St[uid].p)
  (Epoch[uid], Period[uid])  $\leftarrow$  (e', p')
  if expected = 'startedSession' then
    return (e =  $\perp$   $\vee$  e' > e)  $\wedge$  p' = 0
  else if expected = 'joinedSession' then
    return (e =  $\perp$   $\vee$  e' > e)
  else if expected = 'epochChanged' then
    return e' = e + 1  $\wedge$  p' = 0
  else if expected = 'periodChanged' then
    return e' = e  $\wedge$  p' = p + 1
  else if expected = 'eitherChanged' then
    return [(e' = e + 1  $\wedge$  p' = 0)  $\vee$  (e' = e  $\wedge$  p' = p + 1)]
     $\wedge$  e'  $\leq$  Epoch[Leader[uid]]
  return true

Helper: *verifyConsistency(uid, ad)
  (e, p)  $\leftarrow$  (Epoch[uid], Period[uid])
  uidlead  $\leftarrow$  Leader[uid]
  // No assurance after (potential) active attack
  if *triviallyInjectable(uid) then
    return true
  // Group
  if (uid, ad)  $\notin$  *members(uidlead, e, p) then
    return false
  // Keys
  if Key[uidlead, e, p]  $\notin$  { $\perp$ , St[uid].k}  $\vee$  St[uid].k =  $\perp$  then
    return false
  Key[uidlead, e, p]  $\leftarrow$  St[uid].k
  return true

Helper: *setGroup(uidlead, G')
  (e, p)  $\leftarrow$  (Epoch[uidlead], Period[uidlead])
  Group[uidlead, e, p]  $\leftarrow$  G'

Helper: *triviallyInjectable(uid)
  return Corrupted[uid]  $\vee$  Corrupted[Leader[uid]]

Helper: *safe()
  for (uidlead, e, p)  $\in$  Challs do
    if Corrupted[uidlead] then return false
    for uid  $\in$  *members(uidlead, e, p) do
      if Corrupted[uid] then return false
  return true

```

Fig. 11: The helper functions for the cmKEM security game from Fig. 10.

This information is then, among others, used to formalize member authentication in the `*verifyCredentials` helper. In a nutshell, this helper ensures that the protocol only accepts a credential if it is valid. For instance in the `StartSession` oracle, if the protocol accepts (and `try` does not cause the oracle invocation to abort prematurely) then for each participant `uid` we must have `*verifyCredentials(uid)`, or the adversary wins the game.

Note that the definition technically does not define the validity of a concrete credential `sig`, but rather defines the necessary conditions for such a valid credential to exist, which is either when the user is honest, or when the adversary knows the corresponding long-term key `id \in Corrlids`. (Encoding a concrete verification would not improve security, as in both cases the adversary can trivially input a valid credential and we would, hence, only rule out stupid attacks.)

Challenges. The game keeps track of all challenged keys using the `Challs` set. Moreover, to answer challenges consistently it uses the `ChallKeys` array, where once a key for a given state is output it is recorded. (That is, rather than ensuring that only either the `Challenge` or `Test` oracle can be invoked, we simply record the answer of the first such invocation.)

Safety predicate. The `*safe` helper method disallows trivial distinguishing strategies by determining whether the combinations of challenges and corruptions is “safe.” More concretely it checks that if a user `uid` has been corrupted, then no state in which `uid` is a member must have been challenged, as otherwise the adversary could trivially distinguish using the leaked state.

Note that the `*members` helper function tackles a subtlety with respect to being a member and how ratcheting epochs and periods work. Assume that the leader `uidlead` removed a user `uid` while being in epoch `e` and period `p`, causing the leader to advance to the state `e + 1` and period 0. Zoom’s scheme allows however

for a network controlling adversary to instruct another member uid' to instead advance to $(e, p + 1)$ first, for which $\text{Group}[\text{uid}_{\text{lead}}, e, p + 1]$ has never been set. Since uid however could compute the respective key, uid needs to be considered a member of that given state.

Consistency properties. The helpers `*verifyProgress` and `*verifyConsistency` not only update the game’s state but also ensure its consistency properties. First, `*verifyProgress` ensures that a given member ends up in the expected epoch and period. Note that the each call site in the game indicates whether the epoch or period should have been incremented using a special flag. Some special consideration has to be given to the invocation in the `JoinSession` oracle. When uid joins a session for the first time, we allow the party to end up in any arbitrary epoch and period (which `*verifyProgress` will allow, as so far $\text{Epoch}[\text{uid}]$ is not set), while when switching to another session, the party must end up in a strictly greater epoch.

Second, the `*verifyConsistency`(uid, ad) helper first checks that only legitimate members end up in a given state. Note that, for simplicity, we often use the special wildcard symbol `*` for the associated data, i.e., for a given uid and group G , the condition $(\text{uid}, *) \in G$ is true if there exists ad such that $(\text{uid}, \text{ad}) \in G$. Similarly, we write $G \setminus \{(\text{uid}_i, *)\}_{i \in [n]}$ to denote the set $\{(\text{uid}, \text{ad}) \in G : \text{uid} \notin \{\text{uid}_1, \dots, \text{uid}_n\}\}$. Note that, when a participant joins a group we do enforce an exact match of the ad components. Further, `*verifyConsistency` verifies that uid ’s key is the same as their leader’s, i.e., that $\text{Key}[\text{uid}_{\text{lead}}, e, p] = \text{st.k}$ for st denoting uid ’s protocol state. Observe, however, that uid might run ahead of uid_{lead} with respect to the period. To accommodate for this corner case, `*verifyConsistency` actually sets $\text{Key}[\text{uid}_{\text{lead}}, e, p]$ to uid ’s key in case it has not been assigned yet. This ensures that whenever any other member uid' — including uid_{lead} — later moves to the same state, consistency is ensured.

PKI interaction. The `cmKEM` primitive makes use of a PKI for binding a party’s the long-term signing key to their identity id . (For instance, id could be a user name, phone number or email address.) In reality, one has to assume that an adversary can verify any such signature, which is reflected in the game exposing access to `PKI.verify-pk`. Furthermore, the signing key output by the PKI — in Zoom’s protocol — is not exclusively used for the `cmKEM` protocol. We take this into account by exposing an additional signing oracle. Note that, however, that domain separation is ensured by preventing the adversary from using the `cmKEM` context identifier ‘`EncryptionKeyAnnouncement`’.

Advantage. We say that a `cmKEM` scheme is secure if no PPT adversary has a better than negligible advantage in winning the game of Fig. 10. The advantage is defined as

$$\text{Adv}_{\Psi, \mathcal{A}}^{\text{cmKEM}} := 2(\Pr[\text{Sec}_{\Pi, \mathcal{A}}^{\text{LL-CGKA}} \Rightarrow 1] - \frac{1}{2})$$

We intentionally do not define the advantage as a positive quantity by taking its absolute value, as the game returns `true` if the adversary wins by triggering an assertion or correctly guesses the bit b , and `false` if it performs any disallowed operations (such as challenging a key known to a corrupted party).

C.4 Zoom’s `cmKEM` Scheme

Recall Zoom’s `cmKEM` scheme from Section 2.4 and Fig. 1 in particular.

We first establish the following simple result.

Theorem 5. *Zoom’s `cmKEM` scheme, which is presented in Fig. 1, is correct if the underlying nonce-based AEAD and signature schemes are correct.*

Proof. This follows directly from inspection of the protocol. Clearly, if $G = \langle g \rangle$ is a Diffie-Hellman group, then `*encrypt-seed` derives the same symmetric key $k \leftarrow \text{HKDF}(\text{DL}_g(\text{upk}_{\text{uid}_{\text{lead}}}, \text{upk}_{\text{uid}}), \text{‘KeyMeetingSeed’})$ for both the leader uid_{lead} and a respective participant uid . Hence, assuming correctness of the AEAD scheme, uid will decrypt the seed that uid_{lead} actually sent. Let’s first consider the case of a new epoch. Here, uid_{lead} samples a fresh seed' , encrypts this to all participants (including potential new members) and then derives $(\text{seed}, k) \leftarrow \text{PRG.Eval}(\text{seed}')$. The recipients, on the other hand, decrypt seed' and then also apply

(seed, k) \leftarrow PRG.Eval(seed'), obviously resulting in the same key and seed. Second, consider the case of a period change. Here, uid_{lead} encrypts the current seed' (before deriving seed and k) to all fresh parties, who then again derive the correct seed and k , while existing parties are simply told to apply (seed, k) \leftarrow PRG.Eval(seed') as well. \square

Next, we consider the scheme's security.

Theorem 6 (Formal version of Theorem 1). *Zoom's cmKEM scheme presented in Fig. 1 is secure according to the game from Fig. 10, i.e.,*

$$\text{Adv}_{\Psi, \mathcal{A}}^{\text{cmKEM}} = 2(\Pr[\text{Sec}_{\Pi, \mathcal{A}}^{\text{LL-CGKA}} \Rightarrow 1] - \frac{1}{2}) \leq \text{negl}(\kappa),$$

under the Gap-DH assumption, when assuming that Hash is collision resistant, the AEAD scheme is secure, the signature scheme is EUF-CMA secure, the PRG satisfies the standard indistinguishability from random notion, and modeling HKDF as a random oracle.

Proof. We show this proof in three parts. First, we consider a sequence of hybrids leading to a game analogous to $\text{Sec}_{\Psi, \mathcal{A}}^{\text{cmKEM}}$ but with a partly idealized scheme cmKEM' that uses independent shared symmetric key, rather than the actual keys derived from the Diffie-Hellman exchange, for the AEAD. Second, we consider a sequence of hybrids that gradually “disable” winning strategies for the adversary, e.g., where the game the adversary interacts with has certain winning conditions such as **assert** removed. We eventually end up with a game where the only winning strategy for the adversary is to guess the bit b , and show that the adversary's success probability remains unchanged from one such hybrid to the next. Finally, we show that in this “idealized” game guessing the bit b also cannot be done with more than negligible probability.

We first consider a hybrid execution $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$ that behaves like $\text{Sec}_{\Psi, \mathcal{A}}^{\text{cmKEM}}$, but where the adversary wins whenever a collision in Hash (used by Ψ on the associated data) is detected. In particular the game monitors whether, during the execution of Ψ , two different ad and ad' are input such that $\text{Hash}(\text{ad}) = \text{Hash}(\text{ad}')$; if so, then the game immediately returns **true** (in addition to the regular ways of winning the game). Clearly, this hybrid is indistinguishable from the real game by collision resistance of Hash.

Then, we consider another hybrid $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ that behaves like $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$, but uses fresh symmetric keys, independent of the public keys, for the AEAD encryption. More concretely, whenever $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ needs a symmetric key for AEAD.Enc for executing the scheme's ***encrypt-seed** helper — which can happen as part of a **StartSession**, **Add**, or **Remove** oracle invocation — then $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ computes the scheme as follows:

- if the same key — i.e., for the same pair of parties uid and uid' — has been used before, then reuses that key;
- else if either the sender uid or the receiver uid' has been corrupted, or not honestly generated to begin with, then the game computes the proper symmetric key as in Ψ ;
- else it samples a uniform random k of the appropriate length, and stores the tuple $(\text{upk}, \text{upk}', k)$ where upk and upk' are the public key stored in uid and uid' , respectively.

If, later, a user with public key upk'' is corrupted, then the system takes all cached keys $(\text{upk}, \text{upk}', k)$ where either $\text{upk}'' = \text{upk}$ or $\text{upk}'' = \text{upk}'$, computes the respective Diffie-Hellman element (recall that $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ sampled the respective secret key usk''), and programs the random oracle at this position to return k .

Since in $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$ the value k is the output of a random oracle, and thus uniformly randomly distributed, clearly $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ behaves exactly identically unless \mathcal{A} queries the random oracle at one of the positions where it later is programmed. (Or two honest Diffie-Hellman key pairs collide, which only happens with negligible probability.) We now show that triggering such a bad event can be reduced to breaking the Gap-DH assumption.

Claim: Querying the ROM at a position that later needs programming, and thus distinguishing $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$ and $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$, implies breaking Gap-DH.

Proof: Assume N is an upper bound on the number of honest users created in the interaction. Then, the reduction first guesses which of the $N^2/2$ instances (i, j) , the adversary will break. (If the reduction guesses wrong, and the adversary, e.g., corrupts one of those parties, the reduction forfeits. This incurs at most a quadratic loss in the reduction.)

The reduction then emulates $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$ based on the Gap-DH instance as follows: Everything not directly related to the Diffie-Hellman keys is executed just as in $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$. When tasked to honestly create a user, the reduction obtains the two group elements g^a and g^b from the Gap-DH instance and uses them as public keys for parties i and j . (For the other users, the reduction simply samples the secret keys themselves.) As a result, when required to compute a k involving either party, such as $\text{HKDF}(g^{ax}, \text{'KeyMeetingSeed'})$ for a g^x input by the adversary, the reduction may no longer be able to do so. Instead, it proceeds as follows: it first checks that the adversary has not queried the respective group element at the random oracle HKDF — using the DDH oracle — and if such an element is found it uses the k selected by HKDF. Otherwise it just samples a fresh uniform k and for all subsequent random oracle queries checks whether they need to be programmed to this value accordingly. Overall, assuming that we guessed the instance correctly, this reduction thus behaves exactly as $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-1}}$ with the adversary triggering the bad event implying a solution to the Gap-DH instance. \square

This concludes the first part of the proof. So far, we have $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ that is a variant of $\text{Sec}_{\Psi, \mathcal{A}}^{\text{cmKEM}}$ which executes an idealized protocol that uses uniform random symmetric encryption keys for the communication between uid and uid' rather than the ones obtained by applying Diffie-Hellman to the respective embedded public keys upk and upk' , and we have

$$|\Pr[\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}} \Rightarrow 1] - \Pr[\text{Sec}_{\Psi, \mathcal{A}}^{\text{cmKEM}} \Rightarrow 1]| \leq \text{negl}(\kappa) \quad (1)$$

under the Gap-DH assumptions.

Next, we gradually modify the actual security game (rather than the scheme) to disable all its winning conditions except for guessing the bit b .

Claim: Consider the hybrid $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-1}}$ that behaves like $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}}$ except for the helper $\text{*verifyCredentials}(\text{uid})$ always returning **true**. (Analogously, that all statements **assert *verifyCredentials**(uid) are removed.) We claim that those experiments are computationally indistinguishable if the signing scheme is EUF-CMA secure. That is, $|\Pr[\mathcal{H}_{\Psi', \mathcal{A}}^{\text{keys-2}} \Rightarrow 1] - \Pr[\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-1}} \Rightarrow 1]|$ is bounded by the probability of an a PPT \mathcal{A}' breaking the EUF-CMA game, where \mathcal{A}' has roughly the same running time as \mathcal{A} .

Proof: We establish this by observing that in order to have $\text{*verifyCredentials}$ return false in the original game, id must not have been corrupted, i.e., $\text{id} \notin \text{CorrIds}$, and uid must have been recorded as corrupted, i.e., $\text{Corrupted}[\text{uid}] = \text{true}$. The former condition implies that so far $\text{Corrupt}(\text{id})$ has not been invoked (and thus the signing key not leaked to the adversary) whereas the latter condition implies that either uid has not been honestly generated, or uid has been leaked by calling $\text{Corrupt}(\text{id})$. In short, $\text{*verifyCredentials}(\text{uid}) = \text{false}$ implies that $\text{Corrupt}(\text{id})$ has not been invoked and uid has not been honestly generated. Consider now the StartSession oracle. There, $\text{*verifyCredentials}$ is only invoked if the scheme did not reject the given inputs. In the scheme, sig_i is a signature of uid_i with associated data $\text{'EncryptionKeyAnnouncement'}$, which then gets verified as part of the StartSession algorithm, for all $i \in [n]$. Hence, those checks passing for some uid_i without $\text{Corrupt}(\text{id})$ having been invoked or uid_i being the result of an honest CreateUser invocation clearly implies a signature forgery, as such signatures can also not be generated using the game's signing oracle. Analogous arguments can be made for the JoinSession and Add oracles. \square

Claim: Consider a game $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-2}}$ that behaves like $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-1}}$ but has various **assert** statements removed. First, in the CreateUser oracle the following assertions are disabled: the ones about the long-term identity and public key matching the one returned by $\Psi.\text{Identity}$, and the last one about CreateUser returning a valid state (not \perp) and uid being fresh, i.e., $\text{St}[\text{uid}] = \perp$. Second, throughout the game, all assertions with respect to the $\Psi.\text{Meeting}$ algorithm and, finally, in the Add and Remove oracles the respective assertions about the group rosters.

We have that $|\Pr[\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-2}} \Rightarrow 1] - \Pr[\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-1}} \Rightarrow 1]|$ is negligible if the signing scheme is correct.

Proof: In the `CreateUser` oracle, the assertions about `Identity` immediately follow from inspection of the scheme, as does the one about the meeting id and the algorithm returning a valid state. The one about the credentials follows by correctness of the signing scheme, and the `CreateUser` algorithm returning a fresh and unpredictable uid follows from the scheme sampling a fresh and uniform random $\text{upk} \in \mathbb{G}$.

For the remaining assertions related to `Meeting`, observe that they are simply checked beforehand by the respective protocol algorithms that reject any input for which they would be violated. The same holds for the group roster assertions, essentially just reflecting the checks performed by the protocol. Hence, $\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-1}}$ and $\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-2}}$ behave the same assuming the correctness of the underlying signature scheme. \square

Claim: Furthermore, consider $\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-3}}$ that replaces all the `assert *verifyProgress` statements by an invocation of `*verifyProgress` without the adversary winning the game if it returns false. We have that $|\Pr[\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-3}} \Rightarrow 1] - \Pr[\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-2}} \Rightarrow 1]|$ is negligible.

Proof: Again for the most part those checks are trivially satisfied by the way the protocol increments the epoch and period numbers, respectively. Note that if a user has not been in any session before, then `*verifyProgress` in `StartSession` or `JoinSession` always returns true. \square

Claim: Finally, consider $\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-4}}$ that defuses the `*verifyConsistency` assertions, i.e., still invokes the helper methods but without the adversary winning based on the return value. We have that $|\Pr[\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-4}} \Rightarrow 1] - \Pr[\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-3}} \Rightarrow 1]|$ is negligible if the underlying AEAD scheme provides authenticity.

Proof: Observe that `*verifyConsistency` always returns `true` (and the assertion is not triggered) if either uid or their respective current leader uid_{lead} have been corrupted. Hence, in the following we focus on the case where both are honest. First, consider the check of uid being a member of the group failing. That is $(\text{uid}, \text{ad}) \notin \text{*members}(\text{uid}_{\text{lead}}, e, p)$. That in particular means that $(\text{uid}, \text{ad}) \notin \text{Group}[\text{uid}_{\text{lead}}, e, p']$ for all $p' \leq p$. If uid is the leader, then the game enforces that they are always part of the group. Otherwise, the only way for a participant uid to move to this epoch e is by receiving an encryption of a seed for some $p' \leq p$. If both uid and uid_{lead} are honest, and given that no Hash collisions have occurred so far, this implies by the authenticity of the AEAD scheme that either of the two parties must have sent such a message. However, as a participant, uid would not accept such a message from themselves, even if they were a leader at some earlier point, due to enforcing that epochs increment monotonically upon a leader change. Furthermore, authenticity of the AEAD scheme ensures that the participant only accepts the message if the associated data matches. Hence, the only option is for uid_{lead} to actually have sent such message with the matching associated data, in which case (uid, ad) must also be contained in $\text{Group}[\text{uid}_{\text{lead}}, e, p']$.

Next, consider the key consistency check to be failing. Observe that for $p = 0$ all parties (except the leader) only transition upon receiving an encrypted seed, and an honest leader will send the same seed to all participants. Hence, analogous to the argument above, we can conclude that consistency of $\text{Key}[\text{uid}_{\text{lead}}, e, 0]$ is implied by the authenticity of the AEAD scheme, for any uid_{lead} and e. For $p > 0$, observe that the protocol derives those keys deterministically from the previous period. Assume that for uid this check fails, after deterministically hashing forward, for some $p > 0$ — with no previous check having failed at this point. Then, `*verifyConsistency` must have stored a different key k' when verifying another party uid' . Due to the determinism, that means that either those parties uid and uid' had incompatible keys in period $p - 1$, in which case a previous consistency check would have failed, or uid' joined the meeting at this point without having been in $p - 1$. In the latter case, uid' , however, got this key encrypted from uid_{lead} , meaning that uid and uid_{lead} would already have had an inconsistent state at $p - 1$ (if uid had also just joined, they would have received the same seed from uid_{lead} as uid'). Hence, in either case the first check to fail cannot be the key consistency check for a period $p > 0$. \square

In summary, we have $\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-4}}$ that is a variant of $\mathcal{H}_{\psi', \mathcal{A}}^{\text{keys-2}}$ where the winning condition has been modified such that only guessing b counts as winning, and combining the hybrid steps with Eq. (1), we have

$$|\Pr[\mathcal{H}_{\psi', \mathcal{A}}^{\text{checks-4}} \Rightarrow 1] - \Pr[\text{Sec}_{\psi, \mathcal{A}}^{\text{cmKEM}} \Rightarrow 1]| \leq \text{negl}(\kappa) \quad (2)$$

under the given assumptions. We, thus, conclude the proof by showing the following claim.

Claim: We have that guessing b in $\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-4}}$ is hard, i.e.,

$$2(\Pr[\mathcal{H}_{\Psi', \mathcal{A}}^{\text{checks-4}} \Rightarrow 1] - \frac{1}{2}) \leq \text{negl}(\kappa) \quad (3)$$

if the underlying AEAD and the PRG are secure, according to the respective standard notions.

Proof: First, consider an adversary \mathcal{A} that never corrupts a party or injects their own key, i.e., in any interaction `Corrupted[uid] = false` whenever `uid` is used. Now consider a challenge for a state identified by the triple $(\text{uid}_{\text{lead}}, e, p)$. If $p = 0$, then this means that uid_{lead} encrypted a fresh PRG state `seed` to all participants, from which the key k is then derived using $(\text{seed}', k) \leftarrow \text{PRG.Eval}(\text{seed})$. By the security of the AEAD scheme the adversary does not have any information on `seed`, and thus by the PRG security, we can conclude that k is indistinguishable from a uniform random one (as with $b = 1$) that is moreover independent of `seed'`. Now consider the period $p = 1$, in which case the key is derived by simply iterating the PRG once more, i.e., $(\text{seed}'', k') \leftarrow \text{PRG.Eval}(\text{seed}')$. If there are newly added parties in that period, then uid_{lead} encrypts them `seed'`. Again, by AEAD security and security of the PRG, k' is indistinguishable from a uniform random key that is chosen independently of `seed''`. By iterating the same argument for periods $p > 1$, we can conclude that all keys are indistinguishable from independent uniform random ones.

It remains to consider corruptions and convince ourselves that all possible resulting attacks are excluded as trivial by the `*safe` predicate. First, we consider an attacker that now can corrupt parties but will not inject any message to a party `uid` if either `uid` or their current leader uid_{lead} have been corrupted (similar to an honest-but-curious adversary). To this end, assume that a party `uid` either starts out as honest and is then later corrupted as part of a `Corrupt(id)` call, or `uid` was not honestly generated in the first place. Assume that `uid` has been added to the session of a leader uid_{lead} in (e, p) and (potentially) removed in (e', p') . According to `*safe` this means that challenging any state from (e, p) until $e' + 1$ is prohibited. (Observe that `*members` still considers `uid` to be part of $(e', p' + 1)$ and so on.) Clearly, the leader does not send `uid` any information about epochs greater than e' or smaller than e . Hence, it suffices to consider periods before p in epoch e , where we can also observe that the seed `uid` obtains does not reveal information about the previous keys. The same argument extends to cases where either `uid` is part of multiple sessions (that in terms of keys are completely independent) or is later re-added to the same session. Finally, we observe that corrupting a session leader disables all challenges on that given session.

We conclude by considering fully active attacks: Whenever either `uid` or uid_{lead} has been corrupted, then the adversary can make `uid` produce an inconsistent key — with respect to the key the honest leader would produce. We remark, however, that the adversary can only challenge honest keys, since by design `*verifyConsistency` returns early for `uid`, before the injected key could be stored in the game's state. This in particular implies that if the adversary would try to deliver a (mauled) ciphertext from another session to `uid`, either by authenticity of the AEAD scheme `uid` will reject (if `uid` and uid_{lead} are honest) or the adversary is unable to challenge the resulting state, implying that he cannot obtain information from such an injection. In summary, our game simply disallows this kind of (trivial) active attacks. \square

Combining Eqs. (2) and (3) directly yields that $\text{Adv}_{\Psi, \mathcal{A}}^{\text{cmKEM}}$ is negligible, concluding the proof. \square

D Details on LL-CGKA

D.1 Security

In this section, we expand on Section 3.2 and discuss the formal security definition of a LL-CGKA scheme depicted in Figs. 12 and 13.

Game overview. The structure of the game is fairly similar to the one of the `cmKEM` primitive discussed in Section 2.3 and Appendix C.3. It notably differs in the handling of time, additional security guarantees and some other (minor) syntactic differences which we discuss below. (For instance, while the `cmKEM` primitive is a more technical abstraction centered around the notion of a session, the LL-CGKA primitive is tied more stringently to the notion of a meeting and, e.g., replaces `StartSession` with `CatchUp` followed by `Lead`.)

First, observe that the game maintains a *global clock* time that is advanced by the adversary. That is, for each fixed time, the adversary can specify as many operations (such as instructing the leader to add parties or letting a participant process a message) before incrementing time by 1. When creating a user uid , the adversary specifies that parties drift δ with respect to the global clock. This drift is then stored in $Offset[uid]$ and throughout the game used to convert global time to local time whenever any of the LL-CGKA algorithms is invoked; otherwise, all timestamps within the security game refer to global time.

Another important difference is that the LL-CGKA game keeps track of more *past state* in order to formalize stronger properties such as liveness. For instance, while the cmKEM game only kept each user uid 's current epoch and period, the Epoch and Period now track uid 's epoch and period over time. That is, $Epoch[uid, time]$ stores the epoch in which uid was at that *global* time and analogously for $Period[uid, time]$. Moreover, $Leader[uid, e]$ keeps track of uid 's leader for each epoch e . (Recall that each leader change implies an epoch change.)

Key consistency and confidentiality. The game defines the confidentiality and consistency of keys analogous to the cmKEM game. That is, each Challenge query is, depending on the bit b , either answered with the actual protocol key as stored in the Key array, or an independent uniform random one. Similarly, $*verifyConsistency$ ensures that a participant has a consistent view on the key with their respective leader uid_{lead} (and all other participants following the same leader).

The non-triviality condition, ruling out certain combinations of challenges and corruptions, is essentially also still the same: Due to the lack of either FS or PCS, a key is considered to be trivially computable by any member that is or has been part of the respective group, irrespective of when the corruption happens. We stress that, in this regard, the Challenge oracle only allows to challenge keys for which one member has actually been in that state; merely receiving (and setting aside for later) a key without transitioning to the respective epoch e' and period p' does not assign the key to $Key[uid_{lead}, e', p']$.

No front running. In contrast to the cmKEM notion, the LL-CGKA game ensures that all participants only move to a state their leader has already been. The exception is the case where the leader's long-term identity but not their ephemeral identity has been compromised (this is due to in certain settings new states only being authenticated using long-term signing keys). To this end, the game sets keys (and groups) upon each action by the leader. Then, $*verifyConsistency$ ensures that when a participant moves to a new epoch or period, either the leader's long-term identity has been corrupted or the relevant key has already been set.

Progress and credentials. The game ensures that either the epoch or period gets incremented with every change to the group. Since participants are expected to not run ahead of their leader, the difference between epochs and periods, however, becomes mostly irrelevant from a security perspective.¹³ Hence, as a simplification, $*verifyProgress$ simply checks that either of the two has been incremented. Note, however, that for a participant in LL-CGKA not every processed message necessarily triggers an advance to the next group state. As such, for Process the game only checks that the epoch-period pair does not decrease. Progress for participants is then guaranteed as part of liveness (see below).

As in the cmKEM game, $*verifyCredentials$ ensures that the adversary cannot impersonate users without having compromised their long-term key material.

Group roster consistency. The LL-CGKA game formalizes that all parties have a consistent view on the group roster for each state. (Where a state is identified by the leader, epoch, and period respectively.) To this end, $*verifyConsistency$ checks that the group stored as part of a participant uid 's state, $St[uid].G$ matches the intended group by the leader. (Observe to this end, that in all operations changing the group roster, $*setGroup$ stores the intended group of the leader uid_{lead} in $Group[uid_{lead}, e, p]$ *before* $*verifyConsistency$ is invoked for all parties including uid_{lead} .) Note that group roster consistency, in contrast to key consistency, only holds as long as the leader's long-term identity has not been compromised.

Moreover, as a sanity check, $*verifyConsistency$ also ensures that if uid transitions to a state it considers themselves to be part of that state.

¹³ The sole exception is the case where a leader's long-term identity but not ephemeral identity has been compromised.

Main

Procedure: Initialize

```

 $b \leftarrow \{0, 1\}$ 
 $won \leftarrow \text{false}$ 
 $time \leftarrow 1$ 
 $Corrlds, Challs \leftarrow \emptyset$ 
 $St[\cdot], Joined[\cdot], Leader[\cdot, \cdot], Group[\cdot, \cdot, \cdot], Epoch[\cdot, \cdot],$ 
   $Period[\cdot, \cdot], Key[\cdot, \cdot, \cdot], ChallKeys[\cdot, \cdot, \cdot], Offset[\cdot] \leftarrow \perp$ 
 $Corrupted[\cdot] \leftarrow \text{true}$ 

```

Procedure: Finalize(b')

```

if  $\neg *safe()$  then return false
else if won then return true
else return  $b' = b$ 

```

PKI interaction

Oracle: Verify-Pk(id, ipk)
 return PKI.verify-pk(id, ipk)

Oracle: Sign(id, context, m)
 req context $\notin \{\text{'EncryptionKeyAnnouncement', 'LeaderParticipantList'}\}$
 (isk, \cdot) \leftarrow PKI.get-sk(id)
 return Sig.Sign(isk, context, m)

Clock

Oracle: Tick
 $time \leftarrow time + 1$
 $M[\cdot], P[\cdot] \leftarrow \perp$
 for all uid : $St[uid] \neq \perp$ do
 (e, p) \leftarrow (Epoch[uid, time - 1], Period[uid, time - 1])
 (Epoch[uid, time], Period[uid, time]) \leftarrow (e, p)
 if $*isLeader(uid, time)$ then
 (St[uid], M[uid]) \leftarrow
 II.LeaderTick(St[uid], time + Offset[uid])
 assert $*verifyProgress(uid, time, \text{'strict'})$
 if (Epoch[uid, time], Period[uid, time]) \neq (e, p) then
 *setGroup(uid, time, Group[uid, e, p])
 *setKey(uid, time)
 assert $*verifyConsistency(uid, time)$
 P[uid] $\leftarrow *pubState(uid, time)$
 else
 (St[uid], alive, M[uid]) \leftarrow II.ParticipantTick(St[uid], time + Offset[uid])
 if alive then assert $*verifyLiveness(uid, time) \wedge St[uid] \neq \perp$
 else St[uid] $\leftarrow \perp$ // uid drops out
 return (M, P)

User management

Oracle: CreateUser(id, meetingId, offset)
 (ust, uid, sig) \leftarrow II.CreateUser(time + offset, id, meetingId)
 Offset[uid] \leftarrow offset
 assert II.Identity(uid) = id \wedge II.Meeting(uid) = meetingId
 assert St[uid] = \perp
 St[uid] \leftarrow ust
 Joined[uid] \leftarrow time
 Corrupted[uid] \leftarrow false
 return (uid, sig)

Oracle: CatchUp(uid, grpPub)
 req St[uid] $\neq \perp$
 try St[uid] \leftarrow II.CatchUp(St[uid], time + Offset[uid], grpPub)
 return *pubState(uid, time)

Group management

Oracle: Lead(uid_{lead}, {uid_i, sig_i}_i ∈ [n])
 req St[uid_{lead}] $\neq \perp$
 try (St[uid_{lead}], M)
 \leftarrow II.Lead(St[uid_{lead}], time + Offset[uid_{lead}], {uid_i, sig_i}_i ∈ [n])
 for $i \in [n]$ do assert $*verifyCredentials(uid_i)$
 assert $*verifyProgress(uid_{lead}, time, \text{'strict'})$
 *setGroup(uid_{lead}, time, {uid_{lead}, uid₁, ..., uid_n})
 *setKey(uid_{lead}, time)
 assert $*verifyConsistency(uid_{lead}, time)$
 return (M, *pubState(uid_{lead}, time))

Group management (leader)

Oracle: Add(uid_{lead}, {(uid_i, sig_i}_i ∈ [n])
 req St[uid_{lead}] $\neq \perp \wedge *isLeader(uid_{lead}, time)$
 (e, p) \leftarrow (Epoch[uid_{lead}, time], Period[uid_{lead}, time])
 $G \leftarrow$ Group[uid_{lead}, e, p]
 try (St[uid_{lead}], M) \leftarrow II.Add(St[uid_{lead}], time + Offset[uid_{lead}], {(uid_i, sig_i}_i ∈ [n])
 for $i \in [n]$ do
 assert $uid_i \notin G \wedge Meeting(uid_i) = Meeting(uid_{lead})$
 $\wedge *verifyCredentials(uid_i)$
 assert $*verifyProgress(uid_{lead}, time, \text{'strict'})$
 *setGroup(uid_{lead}, time, $G \cup \{uid_i\}_{i \in [n]}$)
 *setKey(uid_{lead}, time)
 assert $*verifyConsistency(uid_{lead}, time)$
 return (M, *pubState(uid_{lead}, time))

Oracle: Remove(uid_{lead}, {uid_i}_i ∈ [n])
 req St[uid_{lead}] $\neq \perp \wedge *isLeader(uid_{lead}, time)$
 (e, p) \leftarrow (Epoch[uid_{lead}, time], Period[uid_{lead}, time])
 $G \leftarrow$ Group[uid_{lead}, e, p]
 try (St[uid_{lead}], M) \leftarrow II.Remove(St[uid_{lead}], time + Offset[uid_{lead}], {uid_i}_i ∈ [n])
 for $i \in [n]$ do
 assert $uid_i \in G \wedge uid_i \neq uid_{lead}$
 assert $*verifyProgress(uid_{lead}, time, \text{'strict'})$
 *setGroup(uid_{lead}, time, $G \setminus \{uid_i\}_{i \in [n]}$)
 *setKey(uid_{lead}, time)
 assert $*verifyConsistency(uid_{lead}, time)$
 return (M, *pubState(uid_{lead}, time))

Participants

Oracle: Follow(uid, m, uid'_{lead}, sig'_{lead})
 req St[uid] $\neq \perp$
 try St[uid] \leftarrow II.Follow(St[uid], time + Offset[uid], m, uid'_{lead}, sig'_{lead})
 assert $*verifyCredentials(uid'_{lead})$
 assert $*verifyProgress(uid, time, \text{'strict'})$
 *setLeader(uid, time, uid'_{lead})
 assert $*verifyConsistency(uid, time)$
 return *pubState(uid, time)

Oracle: Process(uid, m)
 req St[uid] $\neq \perp \wedge \neg *isLeader(uid, time)$
 try St[uid] \leftarrow II.Process(St[uid], time + Offset[uid], m)
 uid_{lead} \leftarrow Leader[uid, Epoch[uid, time]]
 assert $*verifyProgress(uid, time, \text{'weak'})$
 *setLeader(uid, time, uid_{lead})
 assert $*verifyConsistency(uid, time)$
 return *pubState(uid, time)

Challenges & corruptions

Oracle: Test(uid_{lead}, e, p)
 req Key[uid_{lead}, e, p] $\neq \perp$
 if ChallKeys[uid_{lead}, e, p] = \perp then
 ChallKeys[uid_{lead}, e, p] \leftarrow Key[uid_{lead}, e, p]
 return ChallKeys[uid_{lead}, e, p]

Oracle: Challenge(uid_{lead}, e, p)
 req Key[uid_{lead}, e, p] $\neq \perp$
 Challs \leftarrow Challs $\cup \{(uid_{lead}, e, p)\}$
 if ChallKeys[uid_{lead}, e, p] = \perp then
 if $b = 1$ then
 ChallKeys[uid_{lead}, e, p] \leftarrow Key[uid_{lead}, e, p]
 else
 ChallKeys[uid_{lead}, e, p] $\leftarrow \mathcal{K}$
 return ChallKeys[uid_{lead}, e, p]

Oracle: Corrupt(id)
 Corrlds \leftarrow Corrlds $\cup \{id\}$
 $L \leftarrow \emptyset$
 for all uid : St[uid] $\neq \perp$ do
 if Identity(uid) = id then
 Corrupted[uid] \leftarrow true
 $L \leftarrow L \cup \{St[uid]\}$
 (isk, \cdot) \leftarrow PKI.get-sk(id)
 return (L, isk)

Fig. 12: The LL-CGKA security game. We define an LL-CGKA scheme to be secure if no PPT \mathcal{A} can win this game with probability better than negligible above one half.

Game Sec _{\mathcal{I}, \mathcal{A}} LL-CGKA Helpers

```

Helper: *pubState(uid, time)
(e, p) ← (Epoch[uid, time], Period[uid, time])
uidlead ← Leader[uid, e]
return (uidlead, e, p, Group[uidlead, e, p])

Helper: *isLeader(uid, time)
e ← Epoch[uid, time]
return Leader[uid, e] = uid

Helper: *setGroup(uidlead, time, G')
(e, p) ← (Epoch[uidlead, time], Period[uidlead, time])
Group[uidlead, e, p] ← G'
*setLeader(uidlead, time, uidlead)

Helper: *setLeader(uid, time, uidlead)
e ← Epoch[uid, time]
Leader[uid, e] ← uidlead

Helper: *setKey(uidlead, time)
(e, p) ← (Epoch[uidlead, time], Period[uidlead, time])
Key[uidlead, e, p] ← St[uidlead, k][e, p]

Helper: *members(uidlead, e, p)
if Group[uidlead, e, p] ≠ ⊥ then
  return Group[uidlead, e, p]
else if p > 0 then
  return *members(uidlead, e, p - 1)
else return ∅

Helper: *relevantPriorLeaders(uid, time)
// Finds all prior meeting leaders since the
// last time the meeting "started over",
// i.e. the roster changed completely
L ← ∅ // all leaders encountered
G ← {uid} // all users encountered
for e' ← Epoch[uid, time], ..., 0 do
  G' ← ∅ // additional users
  for uid' ∈ G do
    uid'lead ← Leader[uid', e']
    if uid'lead ≠ ⊥ then
      L ← L ∪ {uid'lead}
      for uid'' : Meeting(uid') = Meeting(uid'')
        ∧ uid'lead = Leader[uid'', e'] do
        G' ← G' ∪ {uid''}
  G ← G ∪ G'
return L

Helper: *verifyProgress(uid, time, strictly)
(e, p) ← (Epoch[uid, time], Period[uid, time])
(e', p') ← (St[uid], e, St[uid], p)
(Epoch[uid, time], Period[uid, time]) ← (e', p')
if e ≠ ⊥ then
  if strictly = 'strict' then
    return (e' > e ∨ (e' = e ∧ p' > p))
  else if strictly = 'weak' then
    return (e' > e ∨ (e' = e ∧ p' ≥ p))
return true

Helper: *verifyCredentials(uid)
(id, ·) ← Identity(uid)
return id ∈ CorrIds ∨ ¬Corrupted[uid]

Helper: *verifyConsistency(uid, time)
(e, p) ← (Epoch[uid, time], Period[uid, time])
uidlead ← Leader[uid, e]
// No assurances during (potential) active attack
if *triviallyInjectable(uid) then
  return true
// Keys
if Key[uidlead, e, p] ∉ {⊥, St[uid].k} ∨ St[uid].k = ⊥ then
  return false // basic consistency
if Identity(uidlead) ∉ CorrIds ∧ Key[uidlead, e, p] = ⊥ then
  return false // no front running
Key[uidlead, e, p] ← St[uid].k
// Group
if uid ∉ *members(uidlead, e, p) ∨ uid ∉ St[uid].G then
  return false
if Identity(uidlead) ∉ CorrIds ∧ St[uid].G ≠ Group[uidlead, e, p] then
  return false
// History
if Identity(uidlead) ∉ CorrIds then
  for all uid' : Meeting(uid') = Meeting(uid) ∧ Leader[uid', e] =
  Leader[uid, e] do
    emin ← min(i : Leader[uid, e'] ≠ ⊥ ∧ Leader[uidlead, e'] ≠ ⊥)
    for e' ← emin, ..., e - 1 do
      if Leader[uid, e'] ≠ Leader[uidlead, e'] then
        return false
return true

Helper: *verifyLiveness(uid, time)
(e, p) ← (Epoch[uid, time], Period[uid, time])
uidlead ← Leader[uid, e]
// No assurance after (potential) active attack
for uid'lead ∈ *relevantPriorLeaders(uid, time) do
  if Identity(uid'lead) ∈ CorrIds then
    return true
// Was leader recently here?
Δ ← *liveness-slack(uid, time, Leader, Joined, Offset)
for time' ← time - Δ, ..., time do
  if (e = Epoch[uidlead, time']
    ∧ p ≥ Period[uidlead, time'])
    ∨ e > Epoch[uidlead, time'] then
    return true
return false

Helper: *triviallyInjectable(uid, time)
e ← Epoch[uid, time]
return Corrupted[uid] ∨ Corrupted[Leader[uid, e]]

Helper: *safe()
for (uidlead, e, p) ∈ Challs do
  if Corrupted[uidlead] then return false
  for uid ∈ *members(uidlead, e, p) do
    if Corrupted[uid] then return false
return true

```

Fig. 13: Additional functions used to define the the LL-CGKA security game from Fig. 12.

Meeting history consistency. As an additional property, the game ensures that if any two parties end up in the same state, they share a consistent view on the meeting's history for their common suffix, i.e., since the later one has been added. This, in particular, rules out attacks where a malicious Zoom server instructs different (disjoint) subsets of participants to first follow different leaders, splitting the meeting, and then later merging them under one common leader again.

Due to the key and group roster consistency, enforcing that participants have a consistent view on the history of leaders suffices to ensure an overall consistent view. This is formalized in `*verifyConsistency` by checking that, for each pair of parties that currently have the same leader, epoch, and period, they agree on

the leader of each prior epoch since the last of them joined the meeting. (Observe that enforcing consistency between each uid and their current leader uid_{lead} is insufficient, as uid_{lead} might be much younger than some of the parties.)

Liveness. Finally, we consider our novel liveness property. Liveness is checked as part of the Tick oracle, ensuring that if a participant stays in the meeting after `ParticipantTick`, then their leader must have been in the same state recently, with the adversary winning otherwise. This, in turn, is formalized as part of `*verifyLiveness` by looking for a time within the given liveness slack for which the leader had the same state. Due to the game permitting multiple operations within one clock tick but only recording the latest state for each time, the notion formally cannot check that the leader has been in exactly that state but one that is at least as old. It is easy to see that this implies liveness as well.

Moreover, in the basic security definition, note we do not assure liveness in the presence of active attacks. This is formalized by `*verifyLiveness` checking that none of the leader’s so far could have been (trivially) impersonated by an active attacker. To this end, the `*relevantPriorLeaders` helper computes all the past leaders of a given meeting, except for the case where a meeting at some point completely “started over” (i.e. the leader and set of participants changed completely from one epoch to the next), in which case past leaders no longer matter.

We remark that `*liveness-slack` is a parameter of the security notion, and it is expressed as a function rather than a single value so that the concrete guarantees can depend on the details of the protocol execution, for example how many different leaders a party has encountered. Indeed, Zoom’s current scheme¹ and our improved proposal both achieve this notion with different liveness slack.

D.2 The protocol

In this section, we expand on Section 3.3 by discussing some of the more technical aspects. The protocol is parameterized in a signature scheme `Sig`, a hash function `Hash`, and a `cmKEM` scheme.

Client protocol. Recall the formal description of the client scheme presented in Fig. 3. For simplicity, the scheme is described implicitly maintaining state across invocations — a description of that state can be found in Table 1.

<code>ust.e</code> , <code>ust.p</code> , <code>ust.k[·, ·]</code> , and <code>ust.G</code>	The exported fields as defined in the LL-CGKA syntax.
<code>st</code>	The state of a <code>cmKEM</code> protocol.
<code>me</code>	One’s own ephemeral identity.
<code>uid_{lead}</code>	The identity of the leader.
<code>isk</code>	The user’s (long-term) signing key (<code>ipk</code> denotes the corresponding verification key).
<code>e_{next}</code> and <code>p_{next}</code>	The largest epoch and period, respectively, for which a key has been received (but not necessarily the corresponding LPL).
<code>lplHash</code>	The hash of the last LPL message sent or received.
<code>hbHash</code>	The hash of the last heart-beat message sent or received.
<code>lastHb</code>	An estimate on when the last known heartbeat has been sent by the leader. If none is known, this contains the time when <code>CreateUser</code> was executed.
<code>Added</code> and <code>Removed</code>	Parties added and removed, respectively, since the last time the LPL has been broadcast. (Maintained by the leader only.)
<code>numLplLinks</code>	The number of LPL messages since the latest coalesced one. (For the leader only.)

Table 1: The client’s state of the LL-CGKA scheme.

cmKEM and LPL. Observe that the construction directly leverages a cmKEM scheme for the key exchange. To generate the LPL messages, a leader uid_{lead} maintains the sets of parties added and removed since broadcasting the last LPL. For sending a differential LPL in the `*send-LPL` helper algorithm, uid_{lead} sends out those sets (alongside the other information), while for a coalesced LPL uid_{lead} sends out their current group roster G . Each LPL message further contains a monotonic counter v that is maintained across leader changes.

Heartbeats and liveness. In the following, we point out a number of subtleties with respect to the liveness component of the scheme. The `lastHb` variable denotes the (estimated) *send time* of the last known heartbeat. For a leader uid_{lead} , this simply corresponds to the actual sending time; for a participant uid , however, it corresponds to the last received heartbeat’s indicated timestamp after accounting for any correctable clock drift (see the `*update-liveness` helper). As a result, uid will drop out in `ParticipantTick` whenever $\text{time} > \text{lastHb} + \Delta_{\text{live}}$. We slightly abuse the `lastHb` variable, moreover, to ensure that joining the meeting takes no longer than Δ_{live} . To this end, when creating the ephemeral identity we simply initialize `lastHb` to the current time. While it does not correspond to anything heartbeat related at this point, the regular liveness mechanism then takes care that after Δ_{live} the party is deemed stale and no longer joins.

The protocol keeps track of an upper bound δ on the clock drift with the current leader. To this end, the protocol estimates the drift as the difference between the timestamp time' indicated in the heartbeat message, and the local time time it has been received at – that is, the protocol generally assumes no network delay. If the network does exhibit delays, then the drift estimates $\text{time} - \text{time}'$ exceeds the actual drift, and as such the party overestimates the recentness of the latest heartbeat, which results in the protocol prioritizing correctness over liveness. Note that for subsequent heartbeats from the same leader the drift estimate can only decrease. If the perceived drift becomes larger, the party can conclude that the later heartbeat must have been delayed in transit.

Server protocol. The sever protocol (implicitly) maintains the following state per meeting `meetingId`, which is initialized by the `Init` algorithm:

- The server state of the cmKEM protocol `pubcmKEM[meetingId]`.
- The current group roster $G[\text{meetingId}]$ consisting of the ephemeral identities.
- The list of LPL messages `lpls[meetingId]` since and including the last coalesced one, represented as a FIFO queue.
- The last heartbeat `hb[meetingId]`.

We remark that the server does *not* have to worry about concurrent sessions within a given meeting, as those should never occur with an honest server that orchestrates the execution well. Hence, the server can keep just one state per meeting rather than per meeting and session.

The protocol is depicted in Fig. 14. Once more, we write the `Split` and `GroupState` algorithms to be implicitly stateful with `Init` simply setting up the state (rather than outputting it). The `split` algorithm considers two main cases. If the message M contains a cmKEM message M_K , then it splits this using the respective cmKEM algorithm. This produces a resulting cmKEM share for each current group member that the server then just forwards alongside any potential LPL or heartbeat messages. Additionally, it uses those shares to derive the current group roster. If M does not contain a cmKEM message, then the server just forwards the LPL message and heartbeat to the last known roster.

The `GroupState` algorithm simply outputs the latest cmKEM epoch, the queue of LPL messages, and the last heartbeat. The LPL queue gets updated as part of `Split` using the `*process-LPL` helper algorithm, and contains all LPL messages since the last coalesced one.

D.3 Proof of Theorem 2

In this section, we prove our main result. First, we translate the liveness slack of Theorem 2 into the respective helper function `*liveness-slack` of the formal security game as follows.

Protocol Zoom's Server LL-CGKA

```

Algorithm: Init()
pubcmKEM ← cmKEM.InitSplitState()
G[·] ← ∅
lpls[·] ← ∅ // Empty FIFO-queue
hb[·] ← ⊥

Algorithm: GroupState(meetingId)
grpPub ← (pubcmKEM.E[meetingId], lpls[meetingId], hb[meetingId])
return grpPub

Helper: *process-LPL(meetingId, lpl)
parse (v, coalesced, ...) ← lpl
if coalesced then
  lpls[meetingId] ← ∅ // Restart LPL queue
  lpls[meetingId].enq(lpl)

Algorithm: Split(M)
parse (uidlead, MK, lpl, hb) ← M
meetingId ← cmKEM.Meeting(uidlead)
ms[·] ← ⊥
if MK ≠ ⊥ then
  try (pubcmKEM, msK) ← cmKEM.Split(pubcmKEM, MK)
  G[meetingId] ← ∅
  for all uid : MK[uid] ≠ ⊥ do
    G[meetingId] ← G[meetingId] ∪ {uid}
    ms[uid] ← (MK[uid], lpl, hb)
else
  for all uid : G[meetingId] do
    ms[uid] ← (⊥, lpl, hb)
if hb ≠ ⊥ then
  hb[meetingId] ← hb
if lpl ≠ ⊥ then
  *process-LPL(lpl, meetingId)
return ms

```

Fig. 14: The server protocol of the LL-CGKA scheme.

Definition 3. Consider the following liveness slack algorithm

$$*liveness\text{-}slack(\text{uid}, \text{time}, \text{Leader}, \text{Joined}, \text{Offset}) := \min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]) + \Delta_{\text{live}},$$

where n is the number of distinct leaders encountered so far and Δ_{live} a fixed parameter.

Using this, we now set out to restate our main theorem.

Theorem 7 (Formal version of Theorem 2). The LL-CGKA scheme from Figs. 3 and 14 is secure according to Figs. 12 and 13 with the liveness slack from Definition 3. That is, for any PPT adversary \mathcal{A} , we have

$$\text{Adv}_{\Pi, \mathcal{A}}^{LL\text{-}CGKA} := 2(\Pr[\text{Sec}_{\Pi, \mathcal{A}}^{LL\text{-}CGKA} \Rightarrow 1] - \frac{1}{2}) \leq \text{negl}(\kappa),$$

for the security parameter κ , if the underlying cmKEM scheme is secure, the signature scheme is EUF-CMA secure, and the hash function is collision resistant.

For ease of presentation, we divide the proof into two parts: the first part considering all security properties except liveness, and a second part considering liveness only.

Lemma 3. The advantage of any PPT adversary \mathcal{A} in winning a modified version of $\text{Sec}_{\Pi, \mathcal{A}}^{LL\text{-}CGKA}$ in which *verifyLiveness always returns true, is negligible in the security parameter κ if the underlying primitives are secure.

Proof. In the following we sketch a simple reduction to the cmKEM game, while assuming that the signatures are unforgeable. In brief, the reduction emulates the LL-CGKA game toward the adversary \mathcal{A} by internally using the cmKEM game such that \mathcal{A} winning the former implies the reduction winning the latter. (Or \mathcal{A} having forged a signature or found a hash collision.)

To this end, the reduction internally runs the LL-CGKA scheme Π as follows: for everything cmKEM related it queries the respective oracles of the cmKEM game, while it computes LPL and heartbeat messages itself. To compute the heartbeats' signatures, the reduction uses the signing oracle exposed by the cmKEM game with the context string 'LeaderParticipantList'.

First, we observe that the two games' states line up in that the state of the LL-CGKA game is just a more fine-grained version of cmKEM game's state. For instance, when the cmKEM game store's Epoch[uid], then the LL-CGKA game stores Epoch[uid, time] such that the two values coincide for the current time. (The analogous holds for Period, while Leader[uid, e] corresponds to Leader[uid] for uid's current epoch.) In particular, we can observe that the reduction can maintain all those arrays consistently with exception of the following:

- The state array $\text{St}[\text{uid}]$. Here, the reduction only maintains the LL-CGKA specific state it needs to execute the protocol. Note that the only time it needs the full state (including the cmKEM part) is upon a corruption, in which case it can simply issue the corruption to the underlying cmKEM game.
- The keys array $\text{Key}[\text{uid}_{\text{lead}}, e, p]$, which it does not maintain. For Challenge and Test queries it directly uses the underlying game, while the emulated $\text{*verifyConsistency}$ checks can simply omit the key-consistency check (therefore potentially returning `true` instead of `false`). This is because the key-consistency checks of the LL-CGKA game and the underlying cmKEM game are essentially the same. (In particular, processing a heartbeat or LPL message cannot cause a key inconsistency in the Zoom protocol.) Hence, the very moment those checks were relevant to properly emulate the $\text{*verifyConsistency}$ of the LL-CGKA game, the underlying cmKEM game is won. As a consequence, omitting those checks does not affect \mathcal{A} 's behavior until \mathcal{A} won the cmKEM game — \mathcal{A} has the same winning probability on the underlying cmKEM game in this reduction as they had in a reduction that perfectly emulated $\text{*verifyConsistency}$.

As a direct consequence of the matching state arrays, we observe that the reduction's use of the cmKEM game to run its scheme is not impeded by the latter game's preconditions — i.e., the preconditions of the two games align. For instance, the preconditions \mathcal{A} must satisfy when calling the Add-oracle in the LL-CGKA game are the same as the ones of the cmKEM game.

We now discuss the remaining differences with respect of some of the game's assurances.

***verifyProgress:** In Zoom's protocol the epoch and period is advanced upon receiving a heartbeat and LPL message. However, upon closer inspection, we observe that it is simply advanced to the one stored from the underlying cmKEM protocol. In order to thus violate the predicate, in reality e_{next} and p_{next} would need to violate the assurance. Furthermore, we can observe that the variant of the cmKEM game is strictly stronger (by enforcing whether the epoch or period progressed depending on the caller). This, implies that any \mathcal{A} that violates the LL-CGKA game's condition (by violating it for e_{next} and p_{next}) for sure also wins the cmKEM game.

***verifyConsistency:** Next, we consider the $\text{*verifyConsistency}$ helpers. By inspection we realize that this time some of LL-CGKA game's checks are stronger than the corresponding checks in the cmKEM game: (1) the additional group and key consistency checks and (2) the new history checks. Hence, we need to prove that violating any of those additional conditions violates the security of one of the other primitives.

To this end consider first the additional key consistency key ensuring that, unless the leader's long-term identity has been corrupted, $\text{Key}[\text{uid}_{\text{lead}}, e, p]$ has already been set when a party transitions to epoch e and period p . Here we observe that the scheme only transitions epoch upon receiving a LPL message for said epoch and period as well as a heartbeat message authenticating that LPL message. Since the latter contains a signature under the leader's long-term signing key, having a party accepting such a message without the leader being in that epoch and period consists a signature forgery. Similarly, consider the additional group consistency check ensuring that the participant agrees with the leader on the group roster. Again, this property is only required to hold as long as the leader's long-term identity has not been compromised. Since the group roster is communicated via the LPL messages which are authenticated via signatures as part of the heartbeat messages, violating this property would again require a signature forgery.

Finally, consider the history consistency check and assume that there are users uid and uid' (belonging to the same meeting) for which the assertion fails. Both users must have received at least one heartbeat from uid_{lead} certifying the current epoch e — for now assume that there is one of the heartbeat of uid_{lead} for e that both parties received. Then, since heartbeats form a hash chain that is checked by the parties, by collision resistance this implies that both users agree on the sequence of all heartbeats up to the moment the later user joined. Since each heartbeat, however, uniquely identifies an epoch and leader — and parties only accept epochs for which they know at least heartbeat — this implies the desired consistency. (In the special case where they parties have no shared heartbeat of their current epoch, one party must just have joined and in particular not been in any other epoch, rendering the check trivial.)

**safe*: Finally, consider the **safe* predicate determining trivial wins. The predicate is identical to the one of the cmKEM game. Hence, any win deemed non-trivial for the LL-CGKA game is also deemed non-trivial for the cmKEM game. \square

Liveness. We now show that an adversary \mathcal{A} can also not break the liveness properties formalized as part of the LL-CGKA game with non-negligible probability. To this end, we consider a sequence of additional lemmas establishing intuitive invariants.

Lemma 4. *Consider an execution of the LL-CGKA protocol Π from Figs. 3, 4 and 14 within $\text{Sec}_{\Pi, \mathcal{A}}^{\text{LL-CGKA}}$, with a PPT \mathcal{A} . Then, whenever **verifyLiveness* is not trivially disabled and for each user uid , we have*

$$\delta_{\text{uid}}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \leq \min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}])$$

where n denotes the number of distinct leaders uid encountered so far and $\delta_{\text{uid}}[\text{uid}_{\text{lead}}]$ refers to uid 's protocol state (i.e., their estimates on the drift to uid_{lead}) while $\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]$ refers to the game state (i.e., the actual drift).

Proof. Clearly, the left-hand side of the equation is only updated whenever uid receives a heartbeat. Moreover, between receiving heartbeats, the bound on the right-hand side monotonically increases. Thus, it suffices to consider the invariant right after processing an incoming heartbeat. In the following, assume that the user uid received the last heartbeat at local time $\text{time}_{\text{uid}} = \text{time} + \text{Offset}[\text{uid}]$ and that this heartbeat contained a timestamp, i.e., the sending time, $\text{time}'_{\text{uid}_{\text{lead}}} = \text{time}' + \text{Offset}[\text{uid}_{\text{lead}}]$. Hence, after invoking **update-drift* we have

$$\begin{aligned} & \delta_{\text{uid}}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & \leq \text{time}_{\text{uid}} - \text{time}'_{\text{uid}_{\text{lead}}} - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & = (\text{time}_{\text{uid}} - \text{Offset}[\text{uid}]) - (\text{time}'_{\text{uid}_{\text{lead}}} - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & = \text{time} - \text{time}' \end{aligned}$$

where the inequality follows from the minima computed as part of **update-drift*.

We now first prove the second part of the bound. To this end, observe that in order for uid to accept a heartbeat, this heartbeat must sign over an LPL link that includes — or that links to a prior LPL link that includes — uid 's fresh ephemeral identity (see **receive-LPL*). By causality we can, thus, conclude that the heartbeat has been sent after uid created their identity, i.e., $\text{Joined}[\text{uid}] \leq \text{time}'$ which directly yields

$$\text{time} - \text{time}' \leq \text{time} - \text{Joined}[\text{uid}].$$

Next, we prove the first part of the bound using induction over the invocations of **update-drift*, i.e., we consider one particular invocation and assume that so far the invariant has been maintained. First, consider the case where uid receives their very first heartbeat. We know that due to the liveness mechanism making uid wait at most Δ_{live} to join, receiving this heartbeat must occur at some global time $\text{time} \leq \text{Joined}[\text{uid}] + \Delta_{\text{live}}$. Combining this with the prior established bound directly leads to the following inequality implying the claim:

$$\text{time} - \text{time}' \leq \text{time} - \text{Joined}[\text{uid}] \leq \text{Joined}[\text{uid}] + \Delta_{\text{live}} - \text{Joined}[\text{uid}] \leq \Delta_{\text{live}}.$$

For subsequent heartbeats we distinguish two cases. If the incoming heartbeat is from the same leader as the current one, then we observe that by definition of the protocol $\delta_{\text{uid}}[\text{uid}_{\text{lead}}]$ only gets smaller (or stays) while the $n \cdot \Delta_{\text{live}}$ term remain unchanged. Hence, the invariant is trivially preserved. Finally, consider the case that the heartbeat is from a new leader. Let $\text{uid}''_{\text{lead}}$ denote the previous leader and $\text{time}''_{\text{uid}''_{\text{lead}}} = \text{time}'' + \text{Offset}[\text{uid}''_{\text{lead}}]$ the (local) time they sent their last heartbeat before uid_{lead} took over. Since uid processed this new heartbeat and hasn't dropped out yet, we know that

$$\text{time} + \text{Offset}[\text{uid}] \leq \text{lastHb} + \Delta_{\text{live}}$$

and using the definition of the protocol variable `lastHb` set in `*update-liveness` we get

$$\begin{aligned}
& \text{lastHb} + \Delta_{\text{live}} \\
&= \text{time}_{\text{uid}_{\text{lead}}}'' + \delta_{\text{uid}}[\text{uid}_{\text{lead}}''] + \Delta_{\text{live}} \\
&= \text{time}'' + (\text{Offset}[\text{uid}_{\text{lead}}''] + \delta_{\text{uid}}[\text{uid}_{\text{lead}}'']) + \Delta_{\text{live}} \\
&\leq \text{time}'' + (n - 1) \cdot \Delta_{\text{live}} + \text{Offset}[\text{uid}] + \Delta_{\text{live}} \\
&= \text{time}'' + n \cdot \Delta_{\text{live}} + \text{Offset}[\text{uid}]
\end{aligned}$$

when using the induction hypothesis in the second last step (on the term in parentheses). Combining the two prior bounds we hence obtain

$$\text{time} - \text{time}' \leq \text{time}'' + n \cdot \Delta_{\text{live}} - \text{time}' \leq n \cdot \Delta_{\text{live}}$$

where in the last step we used that $\text{time}'' \leq \text{time}'$ by causality. \square

It remains to show that this implies the desired liveness properties, as expressed in the following lemma.

Lemma 5. *In the following, consider the game that behaves like $\text{Sec}_{\Pi, \mathcal{A}}^{\text{LL-CGKA}}$, with `*liveness-slack` as defined in Definition 3, but where the winning condition is modified to only account for `*verifyLiveness`. Then the advantage of any PPT adversary \mathcal{A} in winning that game is negligible, if the hash function is collision resistant and the signature scheme EUF-CMA secure.*

Proof. Let time denote the current global time, as defined in the game. Assume that uid_{lead} sent (at local time $\text{time}'_{\text{uid}_{\text{lead}}}$) the last heartbeat which uid received so far, and let lastHb be the variable in uid 's state right after that. Inserting the definition of $\text{lastHb} = \text{time}'_{\text{uid}_{\text{lead}}} + \delta_{\text{uid}}[\text{uid}_{\text{lead}}]$ in the following expression yields

$$\begin{aligned}
& \text{lastHb} - \text{time}'_{\text{uid}_{\text{lead}}} - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\
&= \delta_{\text{uid}}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\
&\leq \min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]),
\end{aligned}$$

where the last step directly follows from Lemma 4. Rearranging the terms, we have

$$\text{lastHb} - \text{Offset}[\text{uid}] \leq \text{time}'_{\text{uid}_{\text{lead}}} - \text{Offset}[\text{uid}_{\text{lead}}] + \min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]),$$

where the left hand side denotes the (global) time that uid thinks that this heartbeat certifies, while the right hand side denotes the actual sending time plus the error term.

The extra Δ_{live} term in Definition 3 accounts for uid waiting that long without receiving heartbeats before dropping out. Moreover, we stress that for each party uid , the game $\text{Sec}_{\Pi, \mathcal{A}}^{\text{LL-CGKA}}$ only enforces liveness as long as all leaders uid encountered so far have been honest, satisfying the respective pre-condition of Lemma 4. \square

Together, Lemmas 3 and 5 imply Theorem 7, concluding the proof.

D.4 Correctness

The security game described so far admits protocols in which parties reject all messages and immediately drop out of the meeting. To rule out such trivial protocols, and show that Zoom's protocol does not exhibit this behavior, in the following we discuss some basic correctness conditions.

The correctness game. We formalize correctness using the game depicted in Fig. 15. This game essentially models an honest execution of the LL-CGKA protocol and ensures that parties (1) agree on keys and participants, (2) don't get stuck, i.e., advance whenever they process a message, and importantly (3) do not drop out due to the liveness mechanism.

To model those guarantees, we consider an execution of the protocol that resembles a (simplified) deployment in a client-server setting.

Meeting flow. All communication is routed via the honest server. For instance, to have a new party to join a meeting the adversary can invoke the `CreateUser` oracle. This then generates a fresh `uid` and informs the server, which in turn then informs the current meeting leader to add `uid` to the meeting. This roughly corresponds to how Zoom deploys their protocol; we ignore access policies and advanced features such as locked meetings or the waiting room here. There are several ways in which a party can leave a meeting: A party can deliberately leave (using the `Leave` oracle) which again informs the server of that request. The same oracle, however, also supports a “silent” mode, modeling a party suddenly dropping off such as when losing internet connection. In the latter case the server is not informed. Rather the server may use the `Expel` oracle to instruct the leader to remove a party at any point in time, for instance when suspecting the party to be no longer actually online. The meeting host kicking out a party would also correspond to informing the server to expel the user — albeit this is not actually modeled in our correctness game. Similarly, for simplicity we only allow one participant at a time to be added or removed from the meeting, except at the beginning or during a leader change.

Note that for parties there is a joint `ProcessOrFollow` oracle. Simply put, the oracle will make the party `Process` a message from their current leader and switch leader using `Follow` when receiving a message from a different leader. Again, this roughly models, how one would actually deploy a LL-CGKA protocol in practice where the honest server is assumed to only forward messages from the current leader.

Honest server. Our correctness game keeps additional state for the server beyond the protocol’s one: `ServerLeader` keeps track of the leader of each meeting `meetingId` while `ServerGroup` keeps track of the current group roster from the viewpoint of the server. Additionally, to deliver signatures `sig` and the public group information `grpPub` to the appropriate parties, the server uses `ServerNewUser` to keep track whether a given `uid` is fresh (i.e., has not received a LL-CGKA message yet) and `ServerNewLeader` whether a given leader has been freshly elected (i.e., has not sent a LL-CGKA message yet). Finally, the server keeps track of each party’s most up-to-date signature using `Sigs`.

(Semi-)passive adversary. One of the main differences to the security game is that we assume a (mostly) passive adversary. In particular, we assume the honestly generated messages to be delivered unaltered and in order. (Zoom uses TLS between clients and the server for “control” messages, including all messages related to the key agreement protocol but not the audio and video streams.) In particular, in the correctness game the adversary cannot suppress messages (as the protocol depends on parties receiving all of them to ensure proper functionality). To this end, the game maintains a *message FIFO queue* `UserMsgs[uid]` for each party (modeling down-link communication between the server and the respective user) as well as a server queue `ServerMsgs[meetingId]` per meeting, modeling up-link communication. (For simplicity we assume that if different parties of the same meeting send something to the server, they get delivered perfectly in order.) In the correctness game, rather than providing all the arguments for the oracle calls, the adversary then merely instructs a party to process the next message in their queue. For ease of presentation, we however still maintain separate oracles and instead each oracle checks that it only is invoked if the *message type* of the message in front of the queue matches. For instance, the adversary may only invoke the `Add` oracle, to instruct the leader to add a party, whenever a message of the corresponding type is queued to be delivered to the leader next — the parties to be added as well as their matching signature are then part of the queued messages.

Note how the adversary has access to all the game and the honest parties’ states, and can request any long term secret keys using `Leak-PKI-Keys`. However, we limit the ways in which the adversary can interact with the game: they cannot directly deliver arbitrary messages to parties other than the honest server. The adversary can create corrupted ephemeral identities using `CreateMaliciousUser`. They can also take over an existing participant’s identity to send messages on their behalf using `SendMaliciousSig`, after which we consider that identity corrupted and stop enforcing their behavior or guaranteeing that they do not drop out. Note that the adversary cannot corrupt the meeting leader, or elect a corrupted party as the leader, and therefore we do not need to add parties to `CorrIds`. Moreover, the adversary is limited to send only messages of type ‘sig’ on behalf of malicious parties, as this is the only type of message the honest server would accept from a non-leader. Indeed, a malicious leader could trivially force parties to drop out by sending malformed messages.

While we believe that, for our scheme, correctness is restored whenever a participant starts following an honest leader again without having dropped out, for simplicity this is not reflected in the definition.

Main

Procedure: Initialize

```
won, lost ← false
time ← 1
Corrupted, CorrIds ← ∅
pub pub ←  $\Pi$ .Init()
pub St[·], Leader[·, ·], Group[·, ·, ·], Epoch[·, ·],
  Period[·, ·], Key[·, ·, ·], Sigs[·], Offset[·] ←  $\perp$ 
pub ServerNewUser[·], ServerNewLeader[·] ← false
pub ServerGroup[·] ← ∅
pub ServerLeader[·], LeaderDrop[·] ←  $\perp$ 
pub UserMsgs[·], ServerMsgs[·] ← {}
```

Procedure: Finalize

```
return won  $\wedge$   $\neg$ lost
```

Clock

Oracle: Tick

```
time ← time + 1
for all meetingId : ServerMsgs[meetingId] ≠ {} do
  // Enforce network bound for server
  parse (time', ...) ← ServerMsgs[meetingId].peek()
  if time' +  $\Delta_{\text{network}} < \text{time}$  then
    lost ← true
for all uid : St[uid] ≠  $\perp$   $\wedge$  uid ≠ Corrupted do
  meetingId ←  $\Pi$ .Meeting(uid)
  // Enforce network bound for users
  parse (time', ...) ← UserMsgs[uid].peek()
  if time' +  $\Delta_{\text{network}} < \text{time}$  then
    lost ← true
  // Enforce bound on leader selection
  if LeaderDrop[meetingId] ≠  $\perp$ 
     $\wedge$  LeaderDrop[meetingId] +  $\Delta_{\text{election}} < \text{time}$  then
      lost ← true
  // Make parties tick
  (e, p) ← (Epoch[uid, time - 1], Period[uid, time - 1])
  (Epoch[uid, time], Period[uid, time]) ← (e, p)
  if *isLeader(uid, time) then
    (St[uid], M) ←  $\Pi$ .LeaderTick(St[uid], time + Offset[uid])
    assert St[uid] ≠  $\perp$ 
    assert *verifyProgress(uid, time, 'weak')
    if (Epoch[uid, time], Period[uid, time]) ≠ (e, p) then
      *setGroup(uid, time, Group[uid, e, p])
      *setKey(uid, time)
      assert *verifyConsistency(uid, time)
    if M ≠  $\perp$  then
      m ← (M, Epoch[uid, time], Period[uid, time])
      ServerMsgs[meetingId].enq((time, uid, 'split', m))
  else
    (St[uid], alive, sig) ←  $\Pi$ .ParticipantTick(St[uid], time + Offset[uid])
    if sig ≠  $\perp$  then
      ServerMsgs[meetingId].enq((time, uid, 'sig', sig))
    if uid ∈ ServerGroup[meetingId] then
      assert alive  $\wedge$  St[uid] ≠  $\perp$ 
      assert *verifyLiveness(uid, time)
```

User management

Oracle: CreateUser(id, meetingId, offset)

```
(ust, uid, sig) ←  $\Pi$ .CreateUser(time + offset, id, meetingId)
Offset[uid] ← offset
assert Identity(uid) = id  $\wedge$  Meeting(uid) = meetingId
assert St[uid] =  $\perp$   $\wedge$  uid  $\notin$  Corrupted
St[uid] ← ust
// Require adversary to select leader, if none
if ServerLeader[meetingId] =  $\perp$   $\wedge$  LeaderDrop[meetingId] =  $\perp$  then
  LeaderDrop[meetingId] ← time
  ServerMsgs[meetingId].enq((time, uid, 'created', sig)) // Inform server
```

Oracle: Leave(uid, silent)

```
req St[uid] ≠  $\perp$ 
St[uid] ←  $\perp$ 
meetingId ←  $\Pi$ .Meeting(uid)
if ServerLeader[meetingId] = uid then
  LeaderDrop[meetingId] ← time
if  $\neg$ silent then // inform server
  ServerMsgs[meetingId].enq((time, uid, 'left',  $\perp$ ))
```

Malicious participants

Oracle: CreateMaliciousUser(uid, sig)

```
req St[uid] =  $\perp$   $\wedge$  uid ≠ Corrupted
Corrupted ← Corrupted  $\cup$  {uid}
ServerMsgs[meetingId].enq((time, uid, 'created', sig))
```

Oracle: SendMaliciousSig(uid, sig)

```
meetingId ←  $\Pi$ .Meeting(uid)
req Sigs[uid] ≠  $\perp$   $\wedge$  uid ≠ ServerLeader[meetingId]
Corrupted ← Corrupted  $\cup$  {uid}
ServerMsgs[meetingId].enq((time, uid, 'sig', sig))
```

Oracle: Leak-PKI-Keys(id)

```
(isk, ipk) ← PKI.get-sk(id)
return (isk, ipk)
```

Group management

Oracle: Lead(uid_{lead})

```
req St[uidlead] ≠  $\perp$ 
parse (time', ·, type, m') ← UserMsgs[uidlead].deq()
req type = 'lead'
parse (m, grpPub) ← m'
uids ← {uid |  $\exists$ sig : (uid, sig) ∈ m}
e ← Epoch[uidlead, time]
if Leader[uidlead, e] =  $\perp$  then
  St[uidlead] ←  $\Pi$ .CatchUp(St[uidlead], time + Offset[uidlead], grpPub)
  (ust', M) ←  $\Pi$ .Lead(St[uidlead], time + Offset[uidlead], m)
  if (ust', M) ≠  $\perp$  then
    St[uidlead] ← ust'
    assert *verifyProgress(uidlead, time, 'strict')
    *setGroup(uidlead, time, uids  $\cup$  {uidlead})
    *setKey(uidlead, time)
    assert *verifyConsistency(uidlead, time)
    m ← (M, Epoch[uidlead, time], Period[uidlead, time])
    ServerMsgs[Meeting(uidlead)].enq((time, uidlead, 'split', m))
    LeaderDrop[Meeting(uidlead)] ←  $\perp$ 
  else
    assert uids  $\cap$  Corrupted =  $\emptyset$ 
```

Oracle: Add(uid_{lead})

```
req St[uidlead] ≠  $\perp$   $\wedge$  *isLeader(uidlead, time)
parse (time', ·, type, m) ← UserMsgs[uidlead].deq()
req type = 'add'
uids ← {uid |  $\exists$ sig : (uid, sig) ∈ m}
(e, p) ← (Epoch[uidlead, time], Period[uidlead, time])
G ← Group[uidlead, e, p]
(ust', M) ←  $\Pi$ .Add(St[uidlead], time + Offset[uidlead], m)
if (ust', M) ≠  $\perp$  then
  St[uidlead] ← ust'
  assert *verifyProgress(uidlead, time, 'strict')
  *setGroup(uidlead, time, G  $\cup$  uids)
  *setKey(uidlead, time)
  assert *verifyConsistency(uidlead, time)
  m ← (M, Epoch[uidlead, time], Period[uidlead, time])
  ServerMsgs[Meeting(uidlead)].enq((time, uidlead, 'split', m))
  else
    assert uids  $\cap$  Corrupted =  $\emptyset$ 
```

Oracle: Remove(uid_{lead})

```
req St[uidlead] ≠  $\perp$   $\wedge$  *isLeader(uidlead, time)
parse (time', ·, type, m) ← UserMsgs[uidlead].deq()
req type = 'remove'
(e, p) ← (Epoch[uidlead, time], Period[uidlead, time])
G ← Group[uidlead, e, p]
(St[uidlead], M) ←  $\Pi$ .Remove(St[uidlead], time + Offset[uidlead], m)
assert *verifyProgress(uidlead, time, 'strict')
*setGroup(uidlead, time, G  $\setminus$  m)
*setKey(uidlead, time)
assert *verifyConsistency(uidlead, time)
m ← (M, Epoch[uidlead, time], Period[uidlead, time])
ServerMsgs[Meeting(uidlead)].enq((time, uidlead, 'split', m))
```

<u>Participants</u>	
<p>Oracle: ProcessOrFollow(uid) req St[uid] $\neq \perp \wedge \neg *isLeader(uid, time)$ parse (time', uid', type, m') $\leftarrow UserMsgs[uid].deq()$ req type = 'process' parse (m, sig', grpPub, e', p') $\leftarrow m'$ e $\leftarrow Epoch[uid, time]$ if uid' = Leader[uid, e] then ust' $\leftarrow II.Process(St[uid], time + Offset[uid], m)$ else if Leader[uid, e] = \perp then St[uid] $\leftarrow II.CatchUp(St[uid], time + Offset[uid], grpPub)$ ust' $\leftarrow II.Follow(St[uid], time + Offset[uid], m, uid', sig')$ assert ust' $\neq \perp$ // no error St[uid] $\leftarrow ust'$ assert *verifyProgress(uid, time, 'weak') assert (Epoch[uid, time], Period[uid, time]) = (e', p') // same as leader *setLeader(uid, time, uid') assert *verifyConsistency(uid, time)</p> <p>Server</p> <p>Oracle: Split(meetingId) parse (time', uid', type, m) $\leftarrow ServerMsgs[meetingId].deq()$ req type = 'split' if uid' = ServerLeader[meetingId] then // Ignore old leaders parse (M, e, p) $\leftarrow m$ (pub, ms) $\leftarrow II.Split(pub, M)$ for all (uid, m) $\in ms$ do assert uid $\in ServerGroup[meetingId]$ if ServerNewUser[uid] $\vee ServerNewLeader[meetingId]$ then sig' $\leftarrow Sigs[uid]$ else sig' $\leftarrow \perp$ if ServerNewUser[uid] then grpPub $\leftarrow II.GroupState(pub, meetingId)$ ServerNewUser[uid] $\leftarrow false$ else grpPub $\leftarrow \perp$ UserMsgs[uid].enq((time, uid', 'process', (m, sig', grpPub, e, p))) if ServerNewLeader[meetingId] then ServerNewLeader[meetingId] $\leftarrow false$</p> <p>Oracle: InitiateAdd(meetingId) uid' $\leftarrow ServerLeader[meetingId]$ req uid' $\neq \perp$ parse (time', uid', type, sig) $\leftarrow ServerMsgs[meetingId].deq()$ req type = 'created' Sigs[uid'] $\leftarrow sig$ ServerNewUser[uid'] $\leftarrow true$ ServerGroup[meetingId] $\leftarrow ServerGroup[meetingId] \cup \{uid\}$ m $\leftarrow \{uid', Sigs[uid']\}$ UserMsgs[uid'].enq((time, \perp, 'add', m))</p> <p>Oracle: UpdateSig(meetingId) parse (time', uid', type, sig) $\leftarrow ServerMsgs[meetingId].deq()$ req type = 'sig' Sigs[uid'] $\leftarrow sig$</p>	<p>Oracle: InitiateRemove(meetingId) uid' $\leftarrow ServerLeader[meetingId]$ req uid' $\neq \perp$ parse (time', uid', type, M) $\leftarrow ServerMsgs[meetingId].deq()$ req type = 'left' meetingId $\leftarrow II.Meeting(uid')$ ServerGroup[meetingId] $\leftarrow ServerGroup[meetingId] \setminus \{uid\}$ m $\leftarrow \{uid', Sigs[uid']\}$ UserMsgs[uid'].enq((time, \perp, 'remove', m))</p> <p>Oracle: Expel(uid') // Like Leave but server initiated meetingId $\leftarrow II.Meeting(uid')$ ServerMsgs[meetingId].enq((time, uid', 'left', \perp))</p> <p>Oracle: ElectLeader(uid') meetingId $\leftarrow II.Meeting(uid')$ uid' $\leftarrow ServerLeader[meetingId]$ req uid' $\neq uid_{lead} \wedge St[uid'] \neq \perp \wedge uid' \notin Corrupted$ // Process pending changes (w/o informing ceding leader) while ServerMsgs[meetingId] $\neq \emptyset$ do parse (time', uid', type, m) $\leftarrow ServerMsgs[meetingId].peek()$ if type = 'split' then Split(meetingId) else if type = 'sig' then UpdateSig(meetingId) else if type = 'created' then ServerMsgs[meetingId].deq() Sigs[uid'] $\leftarrow m$ ServerNewUser[uid'] $\leftarrow true$ ServerGroup[meetingId] $\leftarrow ServerGroup[meetingId] \cup \{uid'\}$ else if type = 'left' then ServerMsgs[meetingId].deq() ServerGroup[meetingId] $\leftarrow ServerGroup[meetingId] \setminus \{uid'\}$ // Detect race conditions if UserMsgs[uid'] $\neq \emptyset$ then // Old leader has pending changes to roster // This will cause diverging, inconsistent, state // Thus, treat old leader as removed ServerGroup[meetingId] $\leftarrow ServerGroup[meetingId] \setminus \{uid_{lead}\}$ // Inform new leader req uid' $\in ServerGroup[meetingId]$ m $\leftarrow \{uid, Sigs[uid]\} \cup \{uid' \in ServerGroup[meetingId] \wedge uid \neq uid_{lead}\}$ if ServerNewUser[uid'] then grpPub $\leftarrow II.GroupState(pub, meetingId)$ ServerNewUser[uid'] $\leftarrow false$ else grpPub $\leftarrow \perp$ UserMsgs[uid'].enq((time, \perp, 'lead', (m, grpPub))) ServerLeader[meetingId] $\leftarrow uid'$ ServerNewLeader[meetingId] $\leftarrow true$</p>

Fig. 15: The correctness game for the LL-CGKA notion. The helper algorithms are the same as in Fig. 13. We define an LL-CGKA scheme to be correct if no PPT \mathcal{A} can win this game with better than negligible probability.

Bounded delay network. Given that liveness demands that parties do drop out if the adversary tries to withhold messages, we need to assume bounded network delay for the protocol to work. In the following, let $\Delta_{network}$ denote a bound on the network delay. This is enforced in the game as follows: If at any point in time the any of the message queues contains a message that has been sitting there for too long (w.l.o.g., the first one to be delivered) then the adversary loses the game. That is, we only consider adversaries to be “valid” in the interaction that do not violate the network bound.

Leader election. Similarly, we also bound the time it takes the server to elect a new leader using $\Delta_{election}$. Otherwise, if the server would not elect a leader for a long time, all participants would just drop out. Note that this is measured from the time the leader drops out in the Leave oracle to the time the next leader

receives the instruction to take over. Further, this bound must also be adhered when the old leader *silently* drops out, for instance after suddenly losing internet connection. Hence, Δ_{election} must be chosen such that the server actually can detect and correct such a situation — we discuss the relationships between these various parameters below. Let us provide a few more comments on the leader election process. First, observe that from the viewpoint of the server there is always an active leader, stored as `ServerLeader[meetingId]`, except when a meeting is just about to start. The game, on the other hand, sets `LeaderDrop[meetingId]` to the current timestamp as soon as the old leader drops out, even if the server is not notified. This field is then cleared out as part of `uidlead` whenever a new leader successfully takes over, and the oracle `Tick` enforces that this process takes no longer than Δ_{election} .

Moreover, observe that the `ElectLeader` oracle either allows to just switch the leader or to additionally process an addition or removal message first. The latter mode differs from separately processing said message in not informing the old leader. In the case of an addition, this can be used to make a new participant immediately the new leader (e.g., to support the join-before-host feature of Zoom meetings) whereas the case of a removal for instance models the old leader being the party to be removed. Finally, notice that race conditions may occur when switching over to a new leader. For instance, if the server suspects the old leader to have lost connection it must elect a new one, even if messages from the old leader might later still arrive. In such a case the old leader proceeded to a state which is now incompatible with the one distributed by the new leader, which continued the meeting at the latest state the server was still aware of. As a result, the server treats the old leader as having dropped out if such a race condition occurs.¹⁴

Correctness assurances. The game checks various correctness properties. In particular, the game enforces the following properties not covered by the security notion:

- *No dropping out:* The game enforces that parties do not drop out unless they intentionally leave or are kicked out. This is formalized by checking the respective return value of `ParticipantTick` whenever a party is still supposed to be in the group.
- *No errors:* While the LL-CGKA syntax defines most operations as fallible, the correctness notion ensures that in case of an honest execution algorithms do not reject. More concretely, whenever the leader is instructed to perform an operation, such as adding or removing a party, then they must succeed and produce a LL-CGKA message. Two exceptions apply: Upon being instructed to take over as a leader with a given group the leader is allowed to reject if the group contains corrupted parties. Similarly, upon being instructed to add additional parties, the leader may reject if any of those parties are corrupted. Participants, on the other hand, must never reject a message generated by the honest leader even if other malicious participants are involved — either the leader already rejected or the meeting must progress as expected.
- *Correct state:* When a participant obtains a LL-CGKA message, processing said message must result in the participant transitioning to the same epoch and period as the leader. If the leader after sending a LL-CGKA message is in epoch e and period p , all participants upon processing their respective message share must transition to the same epoch and period. To this end, the correctness game annotates LL-CGKA messages with their respective epoch and period and then enforces the property as part of `ProcessOrFollow`. Note that annotating LL-CGKA messages is done for bookkeeping by the game only, e and p are not actually sent (separately) in a protocol execution.

In addition, the correctness game also enforces various properties that directly mirror the ones from the security game:

- *Consistency:* Analogous to the security game `*verifyConsistency` checks that all parties in the same state, i.e. parties with the same leader, epoch and period, agree on the key, the group roster as well as the meeting’s history. In combination with ensuring that parties do end up in the expected epoch and period, this ensures that parties do agree on all relevant state.

¹⁴ In the game, this is done by directly inspecting the network queues. In a real implementation this would of course need to be implemented by keeping track of group-change requests sent to a leader without receiving a corresponding LL-CGKA message.

- *Progress*: The game assures in `*verifyProgress` that upon each action the leader moves to the next state, i.e., epoch or period.
- *Liveness*: The correctness game also checks liveness, i.e., the fact that parties at all times are in a state for which their leader has been in recently.

Theorem 8. *The LL-CGKA scheme from Figs. 3 and 14 satisfies correctness under the following constraints:*

- (1) $\Delta_{\text{live}} \geq 4 \cdot \Delta_{\text{network}} + \max(\Delta_{\text{LPL}}, \Delta_{\text{election}})$;
- (2) $\Delta_{\text{live}} \geq \Delta_{\text{heartbeat}} + 2 \cdot \Delta_{\text{network}}$;
- (3) $\Delta_{\text{live}} \geq \Delta_{\text{heartbeat}} + \Delta_{\text{election}} + 2 \cdot \Delta_{\text{network}}$;
- (4) $\Delta_{\text{LPL}} \leq \Delta_{\text{heartbeat}}$.

In particular, when modeling a server who considers a leader to have dropped out and elects a new one immediately after $M \geq 1$ missed heartbeats, then we can set

$$\Delta_{\text{election}} = 2 \cdot \Delta_{\text{network}} + M \cdot \Delta_{\text{heartbeat}}$$

(where one Δ_{network} is the time it takes for the server to get heartbeats and the other the time it takes to inform the new leader) and obtain correctness if the parameters satisfy

$$\Delta_{\text{live}} \geq 6 \cdot \Delta_{\text{network}} + (M + 1) \cdot \Delta_{\text{heartbeat}}.$$

Proof (Sketch). We note that the consistency properties follow by inspection of the scheme and correctness of the underlying primitives. In the following, we give a brief argument on the no-dropout property.

First, we observe that condition (1) implies that parties will join a meeting successfully as long as the meeting's current leader remains active long enough to process their request. It takes up to $2 \cdot \Delta_{\text{network}}$ for the server to receive the join request from the new participant `uid` and deliver it to the current leader, which will immediately generate the respective `cmKEM` message for the new party but might take up to Δ_{LPL} for the next refresh of the LPL and heartbeat (in case the last LPL message has just been sent before receiving the request). It then takes up to another $2 \cdot \Delta_{\text{network}}$ to deliver the message from the leader to the joining participant, at which point the participant will delay dropping out by an additional Δ_{live} . If the leader drops out before processing the request, we know that at most Δ_{election} later a new leader will have been elected (and informed). Since the leader drops out no later than $2 \cdot \Delta_{\text{network}}$ after the new participant created their `uid`, in order for `uid` to not be admitted, we know that no later than $2 \cdot \Delta_{\text{network}} + \Delta_{\text{election}}$ after `uid` got created a new leader takes over and immediately generates `cmKEM` message, LPL message, and heartbeat for the current group that does include `uid`. This is then received by `uid` after at most an additional $2 \cdot \Delta_{\text{network}}$. Hence, (1) ensures that `uid` successfully joins even in that case. Finally, if `uid` is elected to be the new leader, this happens no later than $2 \cdot \Delta_{\text{network}} + \Delta_{\text{election}}$ after `uid` got created.

Next, consider a party who has accepted at least a heartbeat in the meeting. We observe that if a leader sends a heartbeat message at *global* time t that is still accepted by a party, then this party will not drop out before global time $t + \Delta_{\text{live}}$, as $\delta[\text{uid}_{\text{lead}}]$ is an upper bound on the drift between a party and their leader, which makes `lastHb` an upper bound on the (local) time when the heartbeat was sent.

Now assume that the party `uid` still accepted a heartbeat sent at global time t . If the adversary does not switch leaders, the same leader will send the following heartbeat at time $t' \leq t + \Delta_{\text{heartbeat}}$ (where the smaller-equal follows from $\Delta_{\text{LPL}} \leq \Delta_{\text{heartbeat}}$). This new heartbeat will be delivered no later than global time $t' + 2 \cdot \Delta_{\text{network}} \leq t + \Delta_{\text{heartbeat}} + 2 \cdot \Delta_{\text{network}} \leq t + \Delta_{\text{live}}$ by condition (2), and therefore it will be accepted as well.

Finally, consider a leader switch. Assume that the old leader sent their last heartbeat at global time t , before leaving at time $t + \Delta_{\text{heartbeat}}$, i.e., right before being supposed to generate the next heartbeat. By assumption, at global time $t' \leq t + \Delta_{\text{heartbeat}} + \Delta_{\text{election}}$ the new leader will have been notified, which immediately generates a `cmKEM` message as well as LPL and heartbeat messages. By the time $t + \Delta_{\text{heartbeat}} + \Delta_{\text{election}} + 2\Delta_{\text{network}}$ participants will have received those messages and, thus, condition (3) ensures that they still did not drop out and will prolong liveness until at least global time $t' + \Delta_{\text{live}}$. \square

E Details on Improved Liveness

E.1 Additional Interaction: Proof of Theorem 3

Recall that our proposal improves the liveness properties twofold. First, the liveness slack no longer degrades in the number of leader changes. Second, liveness now holds even if a *past leader* has been corrupted or malicious as long as the current leader is honest, which is formalized by altering the `*verifyLiveness` helper in the LL-CGKA security game, as detailed in Fig. 16.




Fig. 16: The strengthened liveness properties compared to Fig. 12.

Theorem 9 (Theorem 3 restated). *The modified LL-CGKA scheme from Fig. 5 is secure according to Fig. 12 with the modified `*verifyLiveness` depicted in Fig. 16 and the following liveness slack*

$$*liveness-slack(\text{uid}, \text{time}, \text{Leader}, \text{Joined}, \text{Offset}) := \min(2 \cdot \Delta_{\text{nonce}} + \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]) + \Delta_{\text{live}},$$

*if the underlying cmKEM scheme is secure, the signature scheme is EUF-CMA secure, and the hash function is collision resistant. Additionally, it is secure with the original (weaker) `*verifyLiveness` from Fig. 12 with*

$$*liveness-slack(\text{uid}, \text{time}, \text{Leader}, \text{Joined}, \text{Offset}) := \min(2 \cdot \Delta_{\text{nonce}} + \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}], n \cdot \Delta_{\text{live}}) + \Delta_{\text{live}},$$

which is at most the slack of the current protocol.

It is easy to see that the scheme only modifies the liveness properties, i.e., all other properties such as key confidentiality, key consistency, and authenticity directly translate over from Theorem 7. The existing liveness term

$$*liveness-slack(\text{uid}, \text{time}, \text{Leader}, \text{Joined}, \text{Offset}) := \min(\text{time} - \text{Joined}[\text{uid}], n \cdot \Delta_{\text{live}}) + \Delta_{\text{live}},$$

in case all past leaders have been honest, also carry over from the proof of Lemma 4. Moreover, the modified `*verifyLiveness` depicted in Fig. 16 is strictly stronger — i.e., ensures liveness in a broader set of circumstances — implying that we can focus on that version of `*verifyLiveness` for the additional Δ_{nonce} term. We, thus, observe that it suffices to strengthen Lemma 4 as follows.

Lemma 6 (Adaptation of Lemma 4). *Consider an execution of the modified LL-CGKA scheme from Fig. 5 within $\text{Sec}_{\Pi, A}^{LL-CGKA}$ with the modified *verifyLiveness Fig. 16. Then, whenever *verifyLiveness is not trivially disabled and for each user uid , we have*

$$\begin{aligned} \delta_{\text{uid}}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ \leq \min(2 \cdot \Delta_{\text{nonce}} + \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]) \end{aligned}$$

where $\delta_{\text{uid}}[\text{uid}_{\text{lead}}]$ refers to uid 's protocol state (i.e., their estimates on the drift to uid_{lead}) while $\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]$ refers to the game state (i.e., the actual drift).

Proof. Recall from the proof of Lemma 4 that it suffices to consider the moments right after a party processed a heartbeat message. Moreover, recall that if the user uid received the last heartbeat at local time $\text{time}_{\text{uid}} = \text{time} + \text{Offset}[\text{uid}]$ and this heartbeat contained a timestamp, i.e., the sending time, $\text{time}'_{\text{uid}_{\text{lead}}} = \text{time}' + \text{Offset}[\text{uid}_{\text{lead}}]$, then after invoking *update-drift we have

$$\delta_{\text{uid}}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \leq \text{time} - \text{time}'.$$

First, we observe that the proof of the $\text{time} - \text{Joined}[\text{uid}]$ — i.e., the fact that a party is never off by more than the duration they spent in a meeting — follows analogous to the current protocol. Indeed, we observe that the respective proof didn't rely on leaders being honest at all (and would even hold against a malicious current leader).

We now prove the term $2 \cdot \Delta_{\text{nonce}} + \Delta_{\text{live}}$ on the bound. First, consider the case where uid receives a first heartbeat from a new leader. Since the party will drop out unless they get this heartbeat within Δ_{live} , and this heartbeat must have been generated after the party generated their identity, we can conclude (analogously as in Lemma 4) that the term can be bounded by Δ_{live} . We henceforth only consider the case of a leader switch.

The protocol enforces that upon a leader switch the first heartbeat of the new leader (and the LPL message) is delivered alongside the first cmKEM message. For the nonce included in the cmKEM message being at most that old and unpredictable, we have that the cmKEM message can have been at most $2\Delta_{\text{nonce}}$ old whenever uid at this point in time. Formally, this holds due to the authenticity of the associated data ad , which is set to the nonce, being enforced as part of the $\text{*verifyConsistency}$ condition in the cmKEM security game from Fig. 10. (Crucially, this property holds irrespective whether uid previously interacted with a malicious leader or not.) Since the protocol verifies that the heartbeat message certifies the cmKEM message's epoch and period, and the new leader is assumed to be honest, we can conclude that this heartbeat is indeed the one the new leader generated at the same time as the cmKEM message. Therefore, we conclude that the drift estimate when processing the first heartbeat processed from an honest leader (except when joining the meeting) is at most $2 \cdot \Delta_{\text{nonce}}$ off.

To keep the equations simpler, we bound the error after processing the first heartbeat from a new leader with the sum of the two terms, $2 \cdot \Delta_{\text{nonce}} + \Delta_{\text{live}}$, instead of their maximum.

For subsequent heartbeats of the same leader, the same argument as in the proof Lemma 4 applies to establish that liveness only improves. \square

E.2 Additional Interaction: Correctness

Our modified protocol introduces nonces sent by participants that a new leader has to acknowledge to establish liveness upon a leader change. More concretely, the protocol accepts any of the two most recent nonces. It remains to show that this does not adversely affect correctness as long as the nonce generation interval Δ_{nonce} is chosen to be not too small, i.e., as long as the two most recent nonces stay relevant enough for the information to disseminate.

Theorem 10. *The modified LL-CGKA scheme from Fig. 5, satisfies correctness if $\Delta_{\text{nonce}} \geq 4 \cdot \Delta_{\text{network}}$ and $2 \cdot \Delta_{\text{nonce}} \geq 4 \cdot \Delta_{\text{network}} + \Delta_{\text{LPL}}$ in addition to the constraints of the base protocol. That is, it satisfies correctness according to the game from Fig. 15 if:*

- (1) $\Delta_{\text{live}} \geq 4 \cdot \Delta_{\text{network}} + \max(\Delta_{\text{LPL}}, \Delta_{\text{election}})$;
- (2) $\Delta_{\text{live}} \geq \Delta_{\text{heartbeat}} + 2 \cdot \Delta_{\text{network}}$;
- (3) $\Delta_{\text{live}} \geq \Delta_{\text{heartbeat}} + \Delta_{\text{election}} + 2 \cdot \Delta_{\text{network}}$;
- (4) $\Delta_{\text{LPL}} \leq \Delta_{\text{heartbeat}}$;
- (5) $\Delta_{\text{nonce}} \geq 4 \cdot \Delta_{\text{network}}$;

In particular, when choosing $\Delta_{\text{nonce}} \geq \Delta_{\text{LPL}}$ then constraint (5) implies constraint (6).

Proof (Sketch). We observe that the modified protocol only introduces the additional nonce freshness checks. Hence, it suffices to argue that the above conditions are sufficient for those checks not to fail. First, consider the case of a leader change. We observe that at the time $\text{ElectLeader}(\text{uid}'_{\text{lead}})$ the nonce the server knows for a particular participant uid is at most $\Delta_{\text{nonce}} + \Delta_{\text{network}}$ old, which models the case of the server just missing the next newer nonce that is still in transit. From the time $\text{ElectLeader}(\text{uid}'_{\text{lead}})$ is called, it then takes at most $3 \cdot \Delta_{\text{network}}$ for uid to receive the first message from $\text{uid}'_{\text{lead}}$, with messages sent from the server to $\text{uid}'_{\text{lead}}$, back to the server, and finally to uid . Hence, at the time uid receives the first message from $\text{uid}'_{\text{lead}}$ the included nonce is at most $\Delta_{\text{nonce}} + 4 \cdot \Delta_{\text{network}}$ old. Since uid accepts the two most recent nonces — and hence as long as it is less than $2 \cdot \Delta_{\text{nonce}}$ old — condition (5) implies that uid does not reject the nonce.

Second, consider the case of uid freshly joining the meeting. In that case, uid distributes their nonce as part of their initial message to the server which is then handed to the current leader uid_{lead} at most $2 \cdot \Delta_{\text{network}}$ after the nonce has been created. The leader will immediately sent the cmKEM message back, acknowledging the nonce within time $4 \cdot \Delta_{\text{network}} + \Delta_{\text{LPL}}$. (Note that while it may take up to Δ_{LPL} longer for the LPL and heartbeat to be received, the freshness of the nonce is evaluated at the time the cmKEM message arrives.) Therefore, uid and accepts the nonce as long as $2 \cdot \Delta_{\text{nonce}} \geq 4 \cdot \Delta_{\text{network}}$, which is implied by (5). The case of a leader change during uid joining is covered by the above argument of a regular leader change. \square

E.3 Leveraging Synchronicity: Proof of Theorem 4

Theorem 11 (Theorem 4 restated). *The modified LL-CGKA scheme from Fig. 7 is secure according to Fig. 12 with the following improved liveness slack*

$$\begin{aligned} *liveness\text{-slack}(\text{uid}, \text{time}, \text{Leader}, \text{Joined}, \text{Offset}) \\ := \min(|\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|, n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]) + \Delta_{\text{live}}, \end{aligned}$$

if the underlying cmKEM scheme is secure, the signature scheme is EUF-CMA secure, and the hash function is collision resistant.

Again, it is easy to see that the scheme only modifies the liveness properties, i.e., all other properties such as key confidentiality, key consistency, and authenticity directly translate over from Theorem 7. We, thus, only consider liveness.

We start by establishing the lemma, stating that the protocol maintains proper bounds on the clock drift. In the remainder of this section, we use the following notational convention: For protocol variables we use subscripts to denote the respective party. For example, we use $\text{lastHb}_{\text{uid}}$ to refer to the variable lastHb as maintained by the protocol of party uid . For time related variables we moreover disambiguate local versus global clocks using those subscripts: Global times (as used by the security game) are denoted without subscript — e.g., time refers to the current global time of the security game — whereas $\text{time}_{\text{uid}} = \text{time} + \text{Offset}[\text{uid}]$ refers to the respective local time as used in the protocol by party uid . (Note that uid 's protocol maintains all variables with respect to their local time, i.e., any time related variable with subscript uid can always be understood with respect to uid 's local clock. Variables like elapsed , which are computed by a client but represent an amount of time and not a specific instant relative to their clock, can be added to both local and global times interchangeably.)

Lemma 7. Consider an execution of the LL-CGKA protocol from Fig. 7 within $\text{Sec}_{\Pi, A}^{LL-CGKA}$, with a PPT A . Then, for each user uid that so far only encountered honest leaders, we have

$$\delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \leq \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}] \leq \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}],$$

where the first and last terms refer to uid 's protocol state (i.e., their estimates on the offset to uid_{lead}) while the middle term refers to the game state (i.e., the actual offset).

Proof. We prove this statement using induction over all invocations of *update-drift . To start, observe that the protocol initializes $\delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]$ and $\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]$ to $-\infty$ and $+\infty$, respectively. Hence, the invariant holds initially, and we only need to show that it is preserved by the *update-drift procedure from Fig. 7.

Now, consider the case that uid at global time time receives a heartbeat from an honest leader uid_{lead} with timestamp $\text{time}'_{\text{uid}_{\text{lead}}}$. Assuming unforgeability of signatures, uid_{lead} actually sent that heartbeat at global time $\text{time}' = \text{time}'_{\text{uid}_{\text{lead}}} + \text{Offset}[\text{uid}_{\text{lead}}] \leq \text{time}$. If $\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]$ gets updated in that execution of *update-drift , at local time $\text{time}_{\text{uid}} = \text{time} + \text{Offset}[\text{uid}]$, it gets set to

$$\begin{aligned} \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] &\leftarrow \text{time}_{\text{uid}} - \text{time}'_{\text{uid}_{\text{lead}}} = (\text{time} + \text{Offset}[\text{uid}]) - (\text{time}' + \text{Offset}[\text{uid}_{\text{lead}}]) \\ &= \text{time} - \text{time}' + \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}] \\ &\geq \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}], \end{aligned}$$

implying the second part of the inequality being preserved.

Analogously, if $\delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]$ is updated in *update-drift , it is set to $\text{earliest}_{\text{uid}} - \text{time}'_{\text{uid}_{\text{lead}}}$. Therefore, to show that the lower bound is preserved throughout the execution, it is enough to establish that the following invariant is preserved:

$$\text{earliest}_{\text{uid}} - \text{time}'_{\text{uid}_{\text{lead}}} \leq \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}],$$

where $\text{earliest}_{\text{uid}}$ denotes the value uid computes in *update-drift . If this is uid 's first heartbeat, i.e., $\text{lastHb}_{\text{uid}}^{\min} = \perp$, then we have that $\text{earliest}_{\text{uid}} = \text{lastHb}_{\text{uid}} = \text{Joined}[\text{uid}] + \text{Offset}[\text{uid}]$, i.e., the (local) time at which the ephemeral identity has been created. Since the heartbeat signs over a LPL message that must contain this ephemeral identity we know that $\text{Joined}[\text{uid}] \leq \text{time}'_{\text{uid}_{\text{lead}}} - \text{Offset}[\text{uid}_{\text{lead}}]$, implying the claim. Otherwise, if $\text{lastHb}_{\text{uid}}^{\min} \neq \perp$, consider first the case where the previous heartbeat has been sent by the same (honest) leader uid_{lead} . Then we have that this previous heartbeat contained the timestamp $\text{time}''_{\text{uid}_{\text{lead}}} = \text{time}'_{\text{uid}_{\text{lead}}} - \text{elapsed}'$ (where $\text{elapsed}'$ denotes the value contained in the current heartbeat), and upon receiving it uid set $\text{lastHb}_{\text{uid}}^{\min} = \text{time}''_{\text{uid}_{\text{lead}}} + \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]$. Therefore, when computing *update-drift for the current heartbeat we have that

$$\begin{aligned} \text{earliest}_{\text{uid}} &= \text{lastHb}_{\text{uid}}^{\min} + \text{elapsed}' \\ &= \text{time}''_{\text{uid}_{\text{lead}}} + \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] + \text{elapsed}' \\ &= \text{time}'_{\text{uid}_{\text{lead}}} + \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \\ &\leq \text{time}'_{\text{uid}_{\text{lead}}} + \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}], \end{aligned}$$

where in the last step we used our assumption that in all prior *update-drift invocations the invariant from the lemma statement has been preserved.

Finally, consider the case of the heartbeat being from a new leader, with $\text{uid}''_{\text{lead}}$ being the prior one. We now adapt the above argument. To this end, consider the last heartbeat of said leader sent at local time $\text{time}''_{\text{uid}''_{\text{lead}}}$, before uid_{lead} taking over. The new leader uid_{lead} will set $\text{elapsed}' = \text{time}'_{\text{uid}_{\text{lead}}} - \text{lastHb}_{\text{uid}_{\text{lead}}}^{\max}$. If uid_{lead} has been part of the meeting before, then $\text{lastHb}_{\text{uid}_{\text{lead}}}^{\max}$ has been set by uid_{lead} when receiving the heartbeat with timestamp $\text{time}''_{\text{uid}''_{\text{lead}}}$, for which we have

$$\text{lastHb}_{\text{uid}_{\text{lead}}}^{\max} = \text{time}''_{\text{uid}''_{\text{lead}}} + \delta_{\text{uid}_{\text{lead}}}^{\max}[\text{uid}''_{\text{lead}}] \geq \text{time}''_{\text{uid}''_{\text{lead}}} + \text{Offset}[\text{uid}_{\text{lead}}] - \text{Offset}[\text{uid}''_{\text{lead}}],$$

where we used the first inequality of the lemma. If uid_{lead} has not been part of the meeting before, $\text{lastHb}_{\text{uid}_{\text{lead}}}^{\max}$ has been set by uid_{lead} to the current local time during CatchUp . For uid to accept the new leader's heartbeat,

that leader must have gotten the correct heartbeat as part of `CatchUp`, which includes the old leader's signature (even if uid_{lead} does not verify that signature). Hence, we can conclude that `CatchUp` happened after the previous leader's last heartbeat and hence the same lower bound on $\text{lastHb}_{\text{uid}_{\text{lead}}}^{\max}$ applies. We can derive

$$\text{elapsed}' \leq (\text{time}'_{\text{uid}_{\text{lead}}} - \text{Offset}[\text{uid}_{\text{lead}}]) - (\text{time}''_{\text{uid}''_{\text{lead}}} - \text{Offset}[\text{uid}''_{\text{lead}}]),$$

i.e., that $\text{elapsed}'$ is a lower bound on the actually elapsed time between the two heartbeats, and thus

$$\begin{aligned} \text{earliest}_{\text{uid}} &= \text{lastHb}_{\text{uid}}^{\min} + \text{elapsed}' \\ &= \text{time}''_{\text{uid}''_{\text{lead}}} + \delta_{\text{uid}}^{\min}[\text{uid}''_{\text{lead}}] + \text{elapsed}' \\ &\leq \text{time}'_{\text{uid}_{\text{lead}}} + \delta_{\text{uid}}^{\min}[\text{uid}''_{\text{lead}}] + \text{Offset}[\text{uid}''_{\text{lead}}] - \text{Offset}[\text{uid}_{\text{lead}}] \\ &\leq \text{time}'_{\text{uid}_{\text{lead}}} + \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}], \end{aligned}$$

where in the last step we used our assumption of $\delta_{\text{uid}}^{\min}[\text{uid}''_{\text{lead}}] \leq \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}''_{\text{lead}}]$. \square

Next, we show that the following bound on the bounds difference. (Unsurprisingly, this resulting difference resembles the liveness assurance of Zoom's original protocol.)

Lemma 8. *For any stage of an execution of the LL-CGKA protocol Π from Fig. 3 within $\text{Sec}_{\Pi, A}^{\text{LL-CGKA}}$, we have that*

$$\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] - \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \leq \min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}])$$

whenever `*verifyLiveness` is not trivially disabled and where n is defined as in Theorem 4.

Proof. We again prove this by induction over the invocations of `*update-drift`, i.e., we consider one particular invocation and assume that so far the invariant has been maintained.

First, consider the case where uid receives their very first heartbeat, with timestamp $\text{time}'_{\text{uid}_{\text{lead}}}$. Due to the liveness mechanism making uid wait at most Δ_{live} to join, we know that this must occur at some global time $\text{time} \leq \text{Joined}[\text{uid}] + \Delta_{\text{live}}$, and therefore in this case proving the bound for the second term in the minimum implies it holds for the first. Moreover, we know that $\text{earliest}_{\text{uid}} = \text{Joined}[\text{uid}] + \text{Offset}[\text{uid}]$. Observe that $\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]$ gets set to $\text{time} + \text{Offset}[\text{uid}] - \text{time}'_{\text{uid}_{\text{lead}}}$ and $\delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]$ to $\text{earliest}_{\text{uid}} - \text{time}'_{\text{uid}_{\text{lead}}}$, implying the claim.

Second, if the heartbeat is from a leader from which uid already received a previous one, then the invariant is trivially preserved as $\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]$ only decreases, while $\delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]$ and the terms on the right side only increase.

Third, consider a heartbeat from a new leader uid_{lead} , when the previous leader¹⁵ $\text{uid}''_{\text{lead}}$ sent the last heartbeat at time $\text{time}''_{\text{uid}''_{\text{lead}}}$. Then, `*update-drift` sets the bounds such that

$$\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] - \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] = \text{time} + \text{Offset}[\text{uid}] - \text{earliest}_{\text{uid}} \leq \text{time} + \text{Offset}[\text{uid}] - \text{time}''_{\text{uid}''_{\text{lead}}} - \delta_{\text{uid}}^{\min}[\text{uid}''_{\text{lead}}], \quad (4)$$

by using that $\text{elapsed}' \geq 0$ (which can be deduced by inspection) to bound $\text{earliest}_{\text{uid}}$. Moreover, since uid is still processing this heartbeat and hasn't dropped out already, we know that

$$\text{lastHb}_{\text{uid}} + \Delta_{\text{live}} \geq \text{time} + \text{Offset}[\text{uid}]$$

and we know (the proof is the same as the one that $\delta^* \leq \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]$ in Lemma 9) that

$$\text{lastHb}_{\text{uid}} - \text{time}''_{\text{uid}''_{\text{lead}}} \leq \delta_{\text{uid}}^{\max}[\text{uid}''_{\text{lead}}].$$

Hence, we can conclude by combining the three above inequalities that

$$\delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] - \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \leq \delta_{\text{uid}}^{\max}[\text{uid}''_{\text{lead}}] - \delta_{\text{uid}}^{\min}[\text{uid}''_{\text{lead}}] + \Delta_{\text{live}},$$

¹⁵ Here we assume $\text{uid}''_{\text{lead}} \neq \text{uid}$, but the case where the participant itself was the previous leader can be treated analogously and leads to stronger guarantees as it resets liveness.

and so if $\text{uid}_{\text{lead}}''$ was the $(n-1)$ -th leader by our assumption we obtain the respective bound for the first term of the minimum. For the second term, consider that

$$\text{earliest}_{\text{uid}} \geq \text{Joined}[\text{uid}] + \text{Offset}[\text{uid}]. \quad (5)$$

We have seen above that this holds when processing the very first heartbeat, and the inequality can be extended to further heartbeats by showing that $\text{earliest}_{\text{uid}}$ only increases over time. In particular when processing each heartbeat

$$\text{earliest}_{\text{uid}} \leftarrow \text{elapsed}' + \text{lastHb}^{\min} = \text{elapsed}' + \text{time}' + \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}''] \geq \text{elapsed}' + \text{time}' + \text{earliest}'_{\text{uid}} - \text{time}' \geq \text{earliest}'_{\text{uid}}$$

where $\text{earliest}'_{\text{uid}}$ was the value of the variable before the last heartbeat was processed. Combining Eqs. (4) and (5) gives the desired bound. \square

We now use the above two lemma to bound the error on the estimated timestamps, i.e., the maximal error a participant uid makes when estimating the time a certain heartbeat has been sent.

Lemma 9. *Assume an honest leader uid_{lead} sends a heartbeat at global time time' , i.e., with timestamp $\text{time}' + \text{Offset}[\text{uid}_{\text{lead}}]$, which gets successfully processed by a participant uid at global time time . After executing `*update-liveness`, we have*

$$|\text{time}' - (\text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}])| \leq \min(|\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|, n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]),$$

i.e., the difference between the actual sending time time' and the estimated sending time $\text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}]$ by uid (converted into global time) is bounded by the minimum over the given three terms.

Proof. We first show the inequality holds for the second and third terms of the minimum using Lemma 8. Let

$$\delta^* := \text{lastHb}_{\text{uid}} - (\text{time}' + \text{Offset}[\text{uid}_{\text{lead}}])$$

denote the correction factor applied by uid in `*update-liveness`. Substituting this into the left-hand side yields

$$\begin{aligned} & |\text{time}' - (\text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}])| \\ &= |\text{time}' - (\delta^* + \text{time}' + \text{Offset}[\text{uid}_{\text{lead}}] - \text{Offset}[\text{uid}])| \\ &= |\delta^* + \text{Offset}[\text{uid}_{\text{lead}}] - \text{Offset}[\text{uid}]| \\ &= |\delta^* - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}])|. \end{aligned} \quad (6)$$

By inspection of `*update-liveness` we observe that δ^* can be written as

$$\delta^* = \begin{cases} \max(0, \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]), & \text{if } \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] \geq 0 \\ \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}], & \text{if } \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] < 0 \end{cases} \quad (7)$$

and applying Lemma 7 lets us deduce that $\delta^* \leq \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]$ in either case. Therefore, we obtain

$$\begin{aligned} & \delta^* - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & \leq \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & \leq \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] - \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \\ & \leq \min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]), \end{aligned}$$

where we used Lemma 7 in the second step and Lemma 8 in the last step. Analogously we can write δ^* as

$$\delta^* = \begin{cases} \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}], & \text{if } \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \geq 0 \\ \min(0, \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}]), & \text{if } \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] < 0 \end{cases} \quad (8)$$

yielding $\delta^* \geq \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}]$. As a result, we furthermore obtain

$$\begin{aligned} & \delta^* - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & \geq \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \\ & \geq \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] - \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] \\ & \geq -\min(n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]), \end{aligned}$$

concluding the proof of the second and third term.

Finally, consider the first term of the bound. First, we observe that in case $\delta^* = 0$, the bound is directly implied by Eq. (6).

If $\delta^* < 0$ we have

$$0 > \delta^* \stackrel{(7)}{=} \delta_{\text{uid}}^{\max}[\text{uid}_{\text{lead}}] \stackrel{\text{Lem. 7}}{\geq} \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]$$

and therefore subtracting $(\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}])$ we obtain

$$-(\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) > \delta^* - (\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]) \geq 0$$

which immediately gives the bound by taking the absolute value.

If $\delta^* > 0$ we have

$$0 < \delta^* \stackrel{(8)}{=} \delta_{\text{uid}}^{\min}[\text{uid}_{\text{lead}}] \stackrel{\text{Lem. 7}}{\leq} \text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]$$

from which we can analogously obtain the same bound. \square

It remains to show that this implies the desired liveness properties, as expressed in the following lemma. (Recall that all other properties have already been proven.)

Lemma 10. *In the following, consider the game that behaves like $\text{Sec}_{\Pi, \mathcal{A}}^{\text{LL-CGKA}}$, with **liveness-slack* as defined in Theorem 4, but where the winning condition is modified to only account for **verifyLiveness*. Then the advantage of any PPT adversary \mathcal{A} in winning that game is negligible.*

Proof. Assume that uid_{lead} sent the last heartbeat uid successfully processed at global time time' . Then, it suffices to show

$$\text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}] \leq \text{time}' + \min(|\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|, n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]),$$

where the left hand side denotes the (global) time that uid thinks that this heartbeat certifies, while the right hand side denotes the actual sending time plus the error term. The slack then simply results from uid waiting for Δ_{live} until dropping out. The result then follows directly from Lemma 9. \square

Together, Lemmas 3 and 10 imply Theorem 4, concluding the proof.

E.4 Leveraging Synchronicity: Correctness

We now formalize correctness of the enhanced scheme. As discussed in Section 4, the scheme's improved liveness assurances come at the cost of degraded correctness. In the following, we show that correctness still holds as long as clocks are synchronized, or the drift is small — we conjecture correctness to hold in additional settings such as small network delay and carefully managed leader changes, but do not make any corresponding formal statement. More formally, we show the following modified correctness theorem where conditions (2) and (3) account for the need of synchronized clocks.

Theorem 12. *The modified LL-CGKA scheme from Fig. 7 satisfies correctness under the following constraints:*

$$(1) \Delta_{\text{live}} \geq 4 \cdot \Delta_{\text{network}} + \max(\Delta_{\text{LPL}}, \Delta_{\text{election}});$$

- (2) $\Delta_{\text{live}} \geq \Delta_{\text{heartbeat}} + 2 \cdot \Delta_{\text{network}} + |\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|$;
- (3) $\Delta_{\text{live}} \geq \Delta_{\text{heartbeat}} + \Delta_{\text{election}} + 2 \cdot \Delta_{\text{network}} + |\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|$;
- (4) $\Delta_{\text{LPL}} \leq \Delta_{\text{heartbeat}}$.

The first and last inequality only depend on protocol parameters, while the middle ones bound the clock offset between any meeting participant and any of their meeting leaders over an execution of the protocol. This can be formalized as a restriction on the adversary's behavior in the correctness experiment.

Proof (Sketch). We remark that the only change with respect to the original protocol is the more stringent liveness check. Hence, we focus exclusively on liveness in the following.

First, we observe that with respect to initially joining the meeting the enhanced protocol remains unchanged to the base protocol: The first heartbeat needs to arrive within Δ_{live} from the time the participant's ephemeral identity uid has been created. As a result, condition (1) still ensures that parties will join a meeting successfully.

Next, consider a party who has accepted at least a heartbeat in the meeting. We observe that if a leader uid_{lead} sends a heartbeat message at *global* time time' that is still accepted by a party uid at time time , then uid delay dropping out until global time

$$\text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}] + \Delta_{\text{live}},$$

for $\text{lastHb}_{\text{uid}}$ as computed during `*update-liveness`. Using Lemma 9 we can thus conclude that uid will not drop out before

$$\begin{aligned} & \text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}] + \Delta_{\text{live}} \\ & \geq \text{time}' - \min(|\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|, n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]) + \Delta_{\text{live}}. \end{aligned}$$

If the adversary does not switch leaders, the same leader will send the following heartbeat at global time $\text{time}'' \leq \text{time}' + \Delta_{\text{heartbeat}}$ (where the smaller-equal follows from $\Delta_{\text{LPL}} \leq \Delta_{\text{heartbeat}}$). This new heartbeat will be delivered no later than global time

$$\begin{aligned} & \text{time}'' + 2 \cdot \Delta_{\text{network}} \\ & \leq \text{time}' + \Delta_{\text{heartbeat}} + 2 \cdot \Delta_{\text{network}} \\ & \leq \text{time}' + \Delta_{\text{live}} - \min(|\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|, n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]), \end{aligned}$$

by condition (2), and therefore it will be accepted as well.

Finally, consider a leader switch. Assume that the old leader sent their last heartbeat at global time time' , before leaving at time $\text{time}' + \Delta_{\text{heartbeat}}$, i.e., right before being supposed to generate the next heartbeat. By assumption, at global time $\text{time}'' \leq \text{time}' + \Delta_{\text{heartbeat}} + \Delta_{\text{election}}$ the new leader will have been notified, which immediately generates a `cmKEM` message as well as `LPL` and heartbeat messages. Let time denote the time at which the participant uid obtains those messages, where clearly $\text{time} \leq \text{time}' + \Delta_{\text{heartbeat}} + \Delta_{\text{election}} + 2\Delta_{\text{network}}$. Using condition (3) and Lemma 9 we can conclude

$$\begin{aligned} & \text{time}' + \Delta_{\text{heartbeat}} + \Delta_{\text{election}} + 2\Delta_{\text{network}} \\ & \leq \text{time}' - \min(|\text{Offset}[\text{uid}] - \text{Offset}[\text{uid}_{\text{lead}}]|, n \cdot \Delta_{\text{live}}, \text{time} - \text{Joined}[\text{uid}]) + \Delta_{\text{live}} \\ & \leq \text{lastHb}_{\text{uid}} - \text{Offset}[\text{uid}] + \Delta_{\text{live}}, \end{aligned}$$

where the last term denotes the time uid would drop out of the meeting. Hence, uid will not have dropped out and still accept the message from the new leader. \square